




## EASY DRIVERS

From GPUs to printers, fix your Ubuntu issues



## TAKE PI PHOTOS

Making mobile snaps easy, even in infrared



## BOOST MEMORY

Make more of Linux with faster memory

# LINUX FORMAT

The **#1** open source mag



## BROWSER WARS!

We test the big-name web browsers to see which is best

# POWER UP FEDORA!

Get started with the most powerful distro from the Red Hat experts!

- Geared for Nvidia AI drivers
- New, faster tiling windows
- Secure reproducible builds
- Mobile desktop modes

## PLUS: HOW TO

- Code your own captcha system
- Better organise your ebook library
- Create a custom VPN with Tailscale

# 100

pages of Linux  
tricks, tips  
& more!

## YOUTUBE HACKING

Follow your subs with API access on the Pi

## SERVER RAID!

Get your multi-disk install sysadmin-ready

## SECURED WI-FI

Ensure you're safe, even over public hotspots

LXF December 2024

FUTURE





**THE  
BRAIN  
TUMOUR  
CHARITY**

**A CURE CAN'T WAIT**

**BRAIN TUMOURS  
MOVE FAST.  
WITH YOUR HELP,  
WE CAN TOO!**

We're working to create a future where brain tumours are curable.  
We urgently need your help to accelerate research.

Text DEFEAT5 to 70507 to donate £5, please help us to find a cure.

**[thebraintumourcharity.org](https://thebraintumourcharity.org)**

© The Brain Tumour Charity 2020. Registered Charity in England and Wales  
(1150054) and Scotland (SC045081)

**FR** Registered with  
**FUNDRAISING  
REGULATOR**



# LINUX FORMAT



## » MEET THE TEAM

This issue, we've been installing cool blue Fedora, so what was the last distribution you tried, and why should our readers give it a spin?



### Jonni Bidwell

Well, Fedora does exactly what a distro should do, so that's novel. Beyond that, I think more people should try Arch. Or if you're worried about the cold, Gentoo or even Linux From Scratch. You can keep your study warm with all that exothermic compilation work.



### Michael Reed

It's not an exotic, unheard-of distro, but I fancy giving Arch Linux a try. Maybe I'm a control freak and I like the idea of being able to decide on system components. Valve's renewed commitment to it on the Steam Deck is another factor that's got me interested.



### Nate Drake

I've switched my daily driver from Ubuntu to the community-maintained Ubuntu Cinnamon. The desktop is not too jarring for Windows veterans, plus there's a number of available 'spices'. These are a dazzling array of add-ons c/o Linux Mint that include new themes, applets and extensions.



### Les Pounder

The last distro I tried was Ubuntu 24.10, a boring answer, I know. I wanted to see how the upgrade process went for both desktop and server users. It turns out that the process is pretty smooth, as long as you know which configuration files need to be tweaked.



### Nick Peers

I'm in the process of building my own Raspberry Pi-powered media streamer, so have been looking at KonstaKANG's unofficial port of LineageOS 21 Android TV. This will give me access to Android TV apps and a TV-friendly user interface. And no more streaming over Wi-Fi!

\*Savings are based on the cover price.

## Powered by infinity!



If there's one reason to try Fedora, it's because there's infinity hiding right in the middle of its logo! Possibly not the reason we deserve, but the reason we need! Who wouldn't want to put the power of infinity at the heart of their system? I've previously avoided Fedora as there's no long-term release; it's Red Hat's development distro, where it goes to test the latest hardware features and newest

software developments. The technology it'll ultimately roll into CentOS Stream and its enterprise-powering Red Hat releases.

If innovations interest you, Fedora is the place to be, from Systemd a good while back to Wayland coupled with KDE Plasma 6 now. You know, Red Hat has been pushing for Nvidia to open source its drivers and is championing 'AI' advancements within Fedora, too. That's all besides more functional advancements such as Flatpak distribution, the container-focused Atomic Desktop, and the Silverblue immutable desktop. These are all cutting-edge technologies you can get your hands on and try today. Jonni will be your guide and has been experimenting with them all on his techno boat in the middle of sleepy Bath!

If you're looking to get more from your existing install, we're testing the big-name browsers to see which you should be using, and have a host of projects, from better ebook libraries, infrared photography, creating a personal VPN, and updating your memory to protecting your public Wi-Fi use and seeing how AMD's latest Zen 5 will help you enjoy Linux even more!

*Neil*

**Neil Mohr** Editor  
neil.mohr@futurenet.com






# Contents



## REVIEWS



**TP-Link BE3600** ..... 19  
Is it still Wi-Fi 7, asks **Brandon Hill**, if there's no 6GHz, as he ruminates on this four-antenna router?



**Ubuntu 24.10** ..... 20  
**Nate Drake** brims with excitement at Canonical's latest offering. Will this distro be a plumed Oriole or just plain 'Orrible?

**Kali 2024.3** ..... 22  
**Nate Drake** explores the latest version of the distro beloved by both hackers and *Mr Robot* fans the world over.

**Expirion Linux 5.0.2** ..... 23  
**Nate Drake** casts aside *systemd* in favour of this lightweight Devuan-based distro. Is *SysVinit* the right way to go?



**Core Keeper** ..... 24  
Keep staff in the dark? The Management is ecstatic at the latest innovation from **Christopher Livingston's** crafting stone.

## POWER UP FEDORA!

Not all heroes wear capes. Jonni Bidwell has got a brand new Fedora to discuss and these hat jokes are wearing the brim thin. See page 32!



CREDIT: Magictorch

## ROUNDUP



**Big-name web browsers** ..... 26  
Browsers connect us to the world wide web, and therefore the outside world. **Michael Reed** examines five of the best for the Linux desktop.

## IN DEPTH



**Meditating on Zen 5** ..... 50  
Discover what the new AMD architecture can do for you and your Linux kernel as we take an in-depth look at Zen 5 in the capable hands of **Chris Szewczyk**.

CREDIT: Getty Images/ E+/Mediterranean, AMD



## Pi USER

**Pi news** ..... 41  
Introduced by **Les Pounder**, who takes a look at a hot new Raspberry Pi competitor, plus more SBC titbits.

**Anthias** ..... 42  
Always one to draw attention, **Les Pounder** delves back into the world of digital signage, using his trusty Raspberry Pi.

**Elegoo Mars 5 Ultra** ..... 43  
When they said Mars needs women, this 3D printer is exactly what **Denise Bertacchi** had in mind!



**Hacking with the YouTube API** ..... 44  
Seeking fame and fortune, **Les Pounder** wonders why he only has 586 subscribers on YouTube. Is it his Blackpool accent?




**Capture infrared photos** ..... 46  
**Mike Bedford** unveils an unusual but attractive form of photography courtesy of a Pi NoIR camera and post-processing.

## CODING ACADEMY

**Completing shell commands** ..... 90  
**Ferenc Deák**'s heart is bursting from his chest; it's not a medical emergency but being so close to completing this series!

**Test if you're human** ..... 94  
**David Bolton** hopes to capture your imagination as he explains how to create a PHP captcha to cut down form spams.



## REGULARS AT A GLANCE

**News** ..... 6  
*WordPress* woes; two decades of Ubuntu; uBlock blocked; FLOSS funding; Android gets more Linux; *Winamp* deletes itself; plus industry insider option; and a look at the latest distro releases.

**Kernel watch** ..... 10

**Answers** ..... 11  
**Neil Bothwick** is open-hearted as well as open source, so is more than happy to help readers with problems regarding Secure Boot, an undetected touchpad, resizing partitions, duplicate files, and more.

**Mailserver** ..... 14  
Readers assail **Neil Mohr** with queries and opinions about the open source world.

**Subscriptions** ..... 16  
Grab your monthly dose of Linux and save a massive 50% off the usual price!

**Back issues** ..... 66  
Get hold of previous *Linux Format* editions.

**Overseas subscriptions** ..... 67  
Get *Linux Format* shipped around the globe.

**HotPicks** ..... 83  
**Mayank Sharma** has been a closeted hoarder for years, but instead of junk in his house, he stuffs his machines with apps such as *PDFSam*, *CPU Info*, *Qrcp*, *TLP*, *Recordbox*, *Newsboat*, *Safe Eyes*, *Hurry Curry!*, *HyperPlay*, *Geeqie* and *Rclone*.

**Next month** ..... 98

## TUTORIALS

**TERMINAL: Easy de-archiving** ..... 56  
**Shashank Sharma** considers himself to be a multitasking maestro, and is impressed with utilities that can do the job of many.

**BASICS: Hardware and drivers** ..... 58  
**Nick Peers** discovers how to use your distro and third-party tools to help profile hardware, find missing drivers and more.

**TAILSCALE: Create a mesh VPN** ..... 62  
Craving constant connections, **Matt Holder** discovers how to use Tailscale to create a mesh network between all of his devices.

**PUPPY LINUX: Secure public Wi-Fi** ..... 68  
**Nate Drake** helps you take your life in your hands as you navigate the murky world of public wireless hotspots.

**LIBRUM: Manage your ebook library** ..... 70  
Bookworm **Nate Drake** helps you galvanise your ebook collection and place it in the cloud with *Librum Reader* and *Calibre*.



**UPGRADE: System memory** ..... 74  
Often found forgetting birthdays and dates, **Neil Mohr** wishes he could upgrade his own memory so easily.

## ADMINISTERIA

**Administeria** ..... 80  
A sometimes overlooked skill, **Stuart Burns** explains the basics of preparing and creating pools of RAID storage.

```
lynsid@rhel8:~$ cat /proc/mdstat
Personalities : [raid1] [linear] [multipath] [raid0] [raid6] [raid5]
md127 : active raid1 sdc1[0] sdd1[1]
          9766302720 blocks super 1.2 [2/2] [UU]
          bitmap: 0/73 pages [0KB], 65536KB chunk

md0 : active raid1 sdb2[1]
          123665408 blocks super 1.2 [2/1] [..U]
          bitmap: 1/1 pages [4KB], 65536KB chunk

unused devices: -none-
lynsid@rhel8:~$ sudo mdadm --detail /dev/md0
/dev/md0:
   Version : 1.2
  Creation Time : Fri Sep 13 00:31:06 2024
   Raid Level : raid1
   Array Size : 123665408 (117.94 GiB 126.63 GB)
  Used Dev Size : 123665408 (117.94 GiB 126.63 GB)
   Raid Devices : 2
  Total Devices : 1
 Persistence : Superblock is persistent

 Intent Bitmap : Internal

   Update Time : Fri Sep 27 20:18:51 2024
     State : active, degraded
   Active Devices : 1
   Working Devices : 1
   Failed Devices : 0
```

# Newsdesk

**THIS ISSUE:** WordPress woes » Celebrating two decades of Ubuntu » uBlock blocked » FLOSS funding boost

## LICENSING

# World War WordPress breaks out online!

WordPress founder Matt Mullenweg clashes with hosting provider WP Engine over contributions to open source development.

**T**he popular user hub WPShout once stated, “WordPress is a factory that makes webpages.” This reflects the fact that it was originally released in 2003 as a medium for publishing blogs, but it now supports other sites such as forums and online stores. As of 2023, the software is used by over 43% of the top 10 million websites.

In recent weeks, a furious debate has been sparked in the community between WordPress co-founder Matt Mullenweg and WP Engine, a major hosting provider for WordPress sites.

Mullenweg, who heads up WordPress.org (the open-source project) and WordPress.com (a commercial hosting service), has accused WP Engine of “strip-mining the WordPress ecosystem” without adequately contributing back to the project.

In late September, frustrations boiled over when Mullenweg took centre stage at WordCamp, a WP Engine-sponsored WordPress conference, stating: “The company [WP Engine] is controlled by Silver Lake, a private equity firm with \$102 million in assets under management. Silver Lake doesn’t give a dang about your open source ideals — it just wants return on capital.”

He went on to urge users to “vote with your wallet” and not support WP Engine, which he clearly felt wasn’t supporting the open source project.

In support of his argument, he then wrote a blog post claiming WP Engine only contributes about 40 hours per week to WordPress

development, while his company Automattic puts in around 100 times this amount.

Mullenweg went on to accuse WP Engine of “profiting off the confusion” caused by the company’s branding, claiming that it needs a commercial licence due to “unauthorised” use of the WordPress trademark.

WP Engine hit back shortly after with a cease and desist letter claiming, among other things, that Mullenweg was making “false, harmful and disparaging” statements. The letter went on to say that he’d tried to extort money from the company before the conference, threatening a “scorched earth nuclear approach” if WP Engine didn’t pay Automattic a “very large sum of money”.

This conflict has had serious consequences for WP Engine users, who were temporarily cut off from Wordpress.org servers. This made it much more difficult to install or update themes and plugins.

The long-term implications for open source aren’t clear. While it’s true that WP Engine doesn’t contribute many development hours to WordPress, its business model doesn’t violate the GPL licence used by the software.

Mullenweg also operates his own WordPress hosting platform, so his attempts to have WP Engine pay for a commercial licence or for using the WordPress trademark could just be an attempt to level the playing field.



Matt Mullenweg used this year’s WordCamp 2024 to blast WP Engine for failing to contribute significantly to WordPress development.

### Via Electronic Mail

Holly Hogan, Chief Legal Officer  
Automattic Inc.  
601 29th Street #343  
San Francisco, CA 94110

Re: Automattic’s Actionable Misconduct Directed to WP Engine

Dear Ms. Hogan:

We are litigation counsel for WP Engine and write to address the serious and repeated misconduct Automattic has directed toward WP Engine over the past several days.

Stunningly, Automattic’s CEO Matthew Mullenweg threatened that if WP Engine did not agree to pay Automattic – his for-profit entity – a very large sum of money before his September 20th keynote address at the WordCamp US Convention, he was going to embark on a self-described “scorched earth nuclear approach” toward WP Engine within the WordPress community and beyond. When his outrageous financial demands were not met, Mr. Mullenweg carried out his threat by making repeated false claims disparaging WP Engine to its customers, its ecosystem

WP Engine fired back with a cease and desist letter, claiming Mullenweg tried to extort money from it before the conference.

DISTRO

# Ubuntu turns 20

Canonical and the community celebrate the 40th release of one of the most popular Linux operating systems.



The Ubuntu website contains a detailed timeline of the OS's release history, along with a celebratory video and interviews with the distribution's users.

**R**eaders who can remember October 2004 will know it was a very different time. The number one hit UK song was Eric Prydz's 'Call on Me', and on 20th October that year, developer Canonical issued the very first stable release of Ubuntu.

Although code-named Warty Warthog, the official version number was 4.10. This reflected the year and month of release. Since 2004, with one exception, Ubuntu has followed this cycle of releasing successive versions every six months.

The original release was based on Debian and updates were supported for 18 months. Version 4.10 shipped with the Gnome 2.08 desktop environment and also bundled software such as *Firefox* and *OpenOffice*.

As high-speed broadband was far less readily available 20 years ago, you could request a free physical install CD via Canonical's Shiplt service (the programme was only discontinued in 2011).

In June 2006, Canonical introduced Ubuntu 6.06. This was the first LTS (Long-Term Support) release, offering three years of updates for desktop users. Later LTS versions of Ubuntu extended support to five years. Security and maintenance updates can also now be extended for another five years through Ubuntu Pro subscriptions, then another two after thanks to Legacy Support.

The most recent release of Ubuntu – 24.10 (see p20) or Oracular Oriole – contains some nostalgic elements to reflect the 20th anniversary of the operating system. This includes a special 'warty brown' accent colour, as well as anniversary wallpapers and the original Ubuntu startup sound upon login (this can be disabled from Ubuntu settings).

Head to <https://ubuntu.com/20years> to see all that the Ubuntu operating system has accomplished over the past two decades.

AD BLOCKING

# Chrome Manifest v3 kills uBlock

## Browser disables extension – users seek alternatives.

**U**sers of uBlock Origin in *Chrome* are finding that the browser automatically disables the extension as it's "no longer supported". This is borne out by a message in the Chrome Web Store.

This was all but inevitable after Google's decision to adopt Manifest v3. The official reason was to enhance browser security and performance. However, given that Google's business model relies heavily on ad revenue, its motives may not be entirely noble.

Manifest v3 restricts how extensions can perform updates. Given that ad blockers rely on daily updated block lists of domains, they're unusable in their current form.

Raymond Hill, uBlock Origin's developer, has created a Manifest v3-compliant version

of the extension named uBO Lite. However, Hill acknowledges its limitations, stating on the project GitHub page, "In general, uBOL will be less effective at dealing with websites using anti-content blockers or minimising website breakage." Due to Manifest v3, the extension also doesn't support dynamic URL filtering.

We also discovered the same alert message on the Chrome Web Store when using the open source *Chromium* browser on which it's based.

After a spat with Mozilla, Hill pulled uBOL from the *Firefox* add-ons page. However, the official uBlock Origin remains. In mid-October, he confirmed in a Twitter post: "uBlock Origin has always worked best on *Firefox*, it has capabilities that are not available on *Chromium*-based browsers regardless of MV2/MV3."

OPINION

# WORKING TOGETHER



**Italo Vignoli** is one of the founders of LibreOffice and the Document Foundation.

**“**This year's annual conference of the LibreOffice community was held in Luxembourg, with a different approach from previous years, involving more presentations on open source software. During the opening session, there were speeches by the minister for research, higher education and digitisation, and by representatives of several Luxembourg governmental bodies, who illustrated the country's evolution towards digitisation. Within this process, the role of open source software is crucial, so much so that the creation of the Open Source Program Office was announced.

Overall, the conference lasted two and a half days, from Thursday to Saturday at noon, with a series of parallel tracks starting on Thursday afternoon on both LibreOffice-related topics – new feature development, quality assurance, marketing, localisation, documentation – and cybersecurity and technology in education. To complete the programme, there was a series of workshops on open source software, and a specific one – dedicated to some 20 students – on the basics of LibreOffice development.

Video recordings of several sessions will be available on the conference website: <https://bit.ly/lxf322conf>





## OPINION

## THE BAD OLD DAYS



**Dave Stokes**  
is a FOSS evangelist at Percona.

“Life before open source software was simpler. You made your one choice of vendor. You bought hardware, software, support and training from that one vendor. You would have all the software you needed if you wanted something and could pay for it.

However, you couldn't share software between vendors or even between hardware in your vendor's product line. You were locked in and paid for that privilege.

Then AT&T engineers wrote an OS – UNIX – that could be compiled for multiple hardware types. Suddenly, you could look at the source code. AT&T's licensing was restrictive, so GNU started developing a UNIX workalike. Then a student posted on a mailing list about a free OS: Linux.

Now we had options. We could see the code, modify it, and run it on various hardware platforms. Open source licences became a social and legal contract to promote cooperation, community and harmony.

But many companies are changing their licences to prevent others from making money. This tarnishes the spirit of cooperation, breaking the social contract.

This could cause a domino effect in which other projects seek protection. If enough change, open source will vanish. Sole-source vendors will be deciding how we work once again.”

## FUNDING

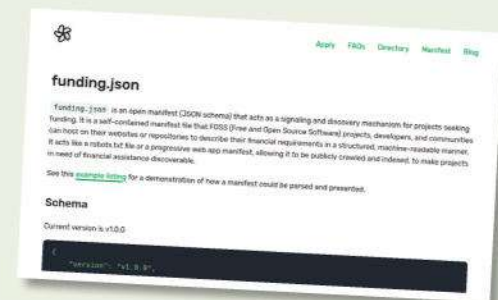
## \$1 million FLOSS fund

Zerodha launches a \$1 million annual fund for open source projects.

**F**inancial services company Zerodha has announced a \$1 million annual fund to support Free/Libre and Open Source Software projects around the world. The initiative, named the FLOSS/fund, will provide \$10,000 to \$100,000 per project annually.

Zerodha CTO Kailash Nadh acknowledged the company had built a commercially successful business “on top of FOSS”, so it was time to “shut up or put up”.

The fund will have a dedicated team to operate in a non-ad-hoc manner. There will also



One of the goals of the funding.json 'experiment' is to provide a way to identify funding requirements of FOSS projects.

be a standardised **funding.json** manifest file, enabling projects to signal their financial needs in a clear and searchable way.

In the official announcement, Nadh also revealed his motivations: “It makes perfect logical sense for a business that relies on FOSS to support it.” At a time when the sustainability of open source has been called into question, the community can only hope that more companies are inspired by Zerodha's example.

See <https://floss.fund/blog/announcing-floss-fund/> for full details.

CREDIT: <https://floss.fund>

## MOBILE

## Android gets more Linux

Here comes the Linux desktop!

**A**ccording to recent Android Open Source Project commits, Google plans to include a native *Terminal* utility in Android 16. This will build on the Android Virtualization Framework, introduced in v13.

The exact implementation isn't clear but ferrochrome-dev-option patches have been submitted to AOSP that include a Linux Terminal toggle under Developer options.

There's no confirmation that *Terminal* will be included in OEM builds of Android but it may feature in Pixel devices.



Android users rely on third-party apps for entering Linux commands, but Google is implementing Terminal into the OS.


## SOFTWARE

## Winamp deletes itself again

Tried open source and didn't like what it found.

**O**n 24th September, *Winamp's* owner the Llama Group released the Legacy Player source code on GitHub. Less than a month later, the repository was removed.

This spectacular volte face was due to a series of blunders. First, the initial Winamp Collaborative License (WCL) forbade coders from forking or modifying the project. It later transpired that the code contained proprietary elements from SHOUTcast server software, Intel and Microsoft. The debacle highlights the challenges of open-sourcing legacy software.



The Winamp legacy code was forked thousands of times before being removed from its GitHub page by the developers.

CREDIT: Google Play Store, GitHub



# Distro watch

What's behind the free software sofa?

## PORTEUX 1.7

This OS shouldn't be confused with Porteus but, like its namesake, is a portable distro based on Slackware Linux. It comes with a wide range of desktop environments (Cinnamon, Gnome, KDE Plasma, LXDE, LXQt, MATE and Xfce) but no preinstalled browser. This must be chosen via the bundled *App Store*. This may be why the latest release uses smaller ISOs than the previous version. There's also a number of optimisations for faster performance. See <https://github.com/porteux/> for more.



PorteuxX is designed to be super-fast, small, portable and modular.

## VOYAGER LIVE 24.10

As the version number suggests, Voyager is based on Ubuntu. However, this French distro adopts a two-to-one approach to the desktop environment by attempting to combine Gnome 47 and Xfce 4.19. Like Ubuntu 24.10, this version will be supported for nine months and uses the Linux 6.11 kernel. It also includes a dedicated gaming profile for Xfce. The operating system has also switched to using *Gnome Software* to handle DEB, Snap and Flatpak packages. Visit <https://voyagerlive.org> to learn more.



French-based distros are our thing!

## SOLUS 4.6

This built-from-scratch distro is available with a variety of desktop environments including Gnome, Xfce (in beta) and Plasma, but the flagship edition deploys Budgie. The default package manager is still *solus-sc* but the latest version continues to have experimental support for *Gnome Software* and *KDE Discover*. The flagship edition also has new default apps. *XReader* has replaced *Evince*, *XViewer* replaces *Eye of Gnome*, and *Enggrampa* replaces *File Roller*. See <https://getsolus/>.



Includes new apps and experimental package manager support.

## CLONEZILLA LIVE 3.2.0-5

Since 2007, this Debian-based live distro has gained a solid reputation for its easy disk management and cloning features. The latest release has upgraded the underlying OS to Debian unstable Sid. The Linux kernel has also been updated to version 6.11.2-1. Certain packages like *reiser4progs* and *wireless-tools* have been removed. Both the *reboot* and *poweroff* commands have been replaced with *systemctl -f reboot*, so root over NFS won't cause the OS to hang when rebooting. Visit <https://github.com/stevenshiau/clonezilla> to learn more.



If only we could clone our brains as easily...

## UBUNTU UNITY 24.10

In 2017, Canonical abandoned its attempt to use its own Unity Desktop on Ubuntu and switched back to Gnome. This didn't sit well with some community members, who continued the OS as a remix. As of 2022, the distro has been an official Ubuntu flavour. The latest release uses Unity 7.7 and, like the previous version, deploys the *Calamares* installer. It also contains bug fixes and switches out *unity-greeter* in favour of *lightdm-gtk-greeter*. Learn more at <https://ubuntuunity.org/>.



The latest version has an overhauled system greeter.

## OPINION

## ARM WRESTLE



**Erik Faye-Lund** is a principal software engineer at Collabora.

Recently, I attended the X.Org Developer's Conference 2024 in Montreal, where I had the opportunity to present, along with my colleague Boris Brezillon, an update on the Panfrost project. What initially began as a reverse engineering effort to understand Arm Mali Midgard and Bifrost GPU internals continues to grow at an amazing pace.

The most recent update involves PanVK, an open source Vulkan driver for those aforementioned GPUs. Support for Arm V10 GPUs landed upstream in Mesa's main branch recently, so it's now possible to start kicking the tyres on Vulkan with an open source driver on Arm Mali-G610 and G310 GPUs.

Note, however, that this is very early support. Neither PanVK itself nor the V10 support is production quality. Additionally, the driver only exposes Vulkan 1.0 with a very minimal set of features, and is in no way conformant yet. As such, don't expect most applications to work out of the box. Some do, but many are likely to fail.

Working closely with Arm, we are working hard on adding features and fixing bugs. For more details, head over to [www.youtube.com/@XOrgFoundation](https://www.youtube.com/@XOrgFoundation) where you'll find the recording of my talk, along with all the other XDC 2024 presentations.

## OPINION

REAL  
TIME

**Jon Masters** is a kernel hacker who's been involved with Linux for over 22 years, and works on energy-efficient Arm servers.

“One of the more spirited discussions in a while began with a patch from Greg Kroah-Hartman titled “MAINTAINERS: Remove some entries due to various compliance requirements”, which systematically removed various maintainers from the kernel's MAINTAINERS file, but had very little context other than saying those removed (who appear largely to have Russian emails or be associated with Russian companies) “can come back in the future if sufficient documentation is provided”. The patch was posted directly to a list intended for raw patches, not the regular LKML (Linux Kernel Mailing List), and was merged promptly without (obvious) public discussion.

It did, however, provoke a discussion on LKML. Linus Torvalds finally weighed in by replying to a post, stating, “OK, lots of Russian trolls out and about ... It's entirely clear why the change was done.” Others later noted apparent legal guidance they'd received on the topic.

Needless to say, things became quite heated. It's too early to say what the long-term impact of this all may be. Perhaps (probably) nothing. One thing is very clear: the kernel community is not immune from facing the impact of the real world.

# Kernel Watch

**Jon Masters** summarises the latest happenings in the Linux kernel, so that you don't have to.

**L**inus Torvalds released Linux 6.12-rc4 (Release Candidate 4), saying that he was “not happy with how big this [the number of changes going in at this stage in the development cycle] is – it's probably far from the biggest rc4 ever, but it is the biggest rc4 we've had in the 6.x series”. Nonetheless, there were no deal-

thousands of simultaneously executing programs by timeslicing these on to the available cores, and handling blocking while programs wait for IO, page faults (loading data in from swap), and so on. Schedulers are very much key to user experience and impact everything from web server response to video game FPS. They must trade off latency and throughput quite carefully.

You may recall that the last of the Real Time (officially PREEMPT-RT) patches were finally merged into the kernel last month, meaning that Linux can now support very low latencies for tasks out of the box.

## À LA MODE

“Linux has a number of different modes of operation now that affect preemption”

breakers and things are on track for a Linux 6.12 final within the next month. Hopefully, this means we will have some early coverage of the Linux 6.13 merge window and its new features in time for next month.

### Lazy preemption muck

Peter Zijlstra is a well-known Linux kernel developer who touches many parts of the kernel, in particular the scheduler. The scheduler decides which task (process) to run at any moment. It is responsible for maintaining the illusion of hundreds or

Technically, Linux has a number of different modes of operation now that affect ‘preemption’, or bumping a task because another one is more important at this moment. Peter's latest work adds another mode called Lazy that aims to ultimately rationalise several of the existing modes. When all is said and done, the goal is to have the option for Real Time, where latency is minimised at a cost of some loss of throughput, and a Lazy mode where latency is pretty good but the kernel won't preempt itself in most cases if another task could run right away. **LXF**

## » ONGOING DEVELOPMENT

Patches continue to be posted for assorted Qualcomm X-Elite Arm platforms, including most recently the (now discontinued) Snapdragon Devkit for Windows, and the Dell XPS 13 model (which this author is selfishly quite interested in testing – we'll report back once there's a chance to see how this works out).

Mario Limonciello posted updated patches for an AMD feature known as the Hardware Feedback Interface (HFI), which apparently provides hardware hints about placement of tasks in heterogeneous core environments. In other words, it appears similar to Intel's Thread Director technology that (on recent Intel chips) provides such placement hints for the P (Performance)

and E (Efficiency) cores in order to ensure that tasks don't get stranded on the wrong type of core.

The RISC-V community is working on updated patches to support its new IOMMU architecture. An IOMMU is used to isolate various hardware components (for example, PCIe devices) and prevent them from being able to unilaterally splat over random memory. Nearly all modern servers (and even PCs) have IOMMUs, because they are necessary to prevent so-called Evil Maid-type attacks, in which malicious or compromised hardware attached to a machine is used to read (or modify) its memory, to extract data or to change the code running within the OS itself.



# Answers



**Neil Bothwick**  
is open-hearted as well as open source.

Got a burning question about open source or the kernel? Whatever your level, email it to [answers@linuxformat.com](mailto:answers@linuxformat.com)

## Q Network booting

I have a Dell Inspiron 3583 laptop that came with Windows 10. I didn't like it and switched to Ubuntu. All went well until I installed 23.10, when my laptop crashed. When I tried to reinstall, all I got was PXE over IPv4. Fortunately, I could run an OS as a live system. I tried several using Ventoy and then tried to install ones that I liked to the internal drive. All the live distros worked but attempts to install just came back with the PXE message. Then I tried Debian 12.5 and it installed without any problem. I tried other distros again but they back came with the PXE message. Of course, I now use Debian 12.5 and I like it, but what is going on?

John Goldspink

A We have seen this before when a system crashes and reverts to a default BIOS setting. That means that Secure Boot is enabled, so the computer will not boot from an OS that is not considered secure. PXE is a network booting protocol; when no suitable hard drive is found to boot from, the firmware looks for a network boot server, fails to find one and gives you that error message, as the last resort for booting has failed.

The reason that Debian worked for you is that it has support for Secure Boot, which requires signed code for the initial boot. Many other distros do not have the signed boot shims needed for Secure Boot, but the solution is remarkably simple. Provided you do not need Secure Boot – and if you have got rid of Windows,

you do not – all you need to do is disable Secure Boot in the BIOS settings and all will be well. Just remember that if you have another hard crash, you may need to repeat this action.

Disabling network boot in the BIOS will stop the error message, but not the lack of booting.

## Q Untouchable laptop

I have an HP laptop and am currently dual booting Ubuntu 24.04 with Windows 11. When I select Ubuntu from the GRUB menu, my Elantech touchpad is not detected. This is not due to my recent upgrade to 24.04 but has been an ongoing issue for years. I have noted on my, always fruitless, search for a solution to this problem that it is suffered by many other laptop users and covers many other laptop manufacturers. I have seen requests from so many who have posted on a variety of forums going as far back to Ubuntu 14.04 at least, and also for many other distros, for a solution to this self same problem. The following is my only, if ugly, method of getting a working touchpad on my laptop: when my GRUB menu first appears, I select the UEFI Settings option. I then select the Boot Options from the next screen, which opens the Boot Loader showing Ubuntu and Windows. On selecting Ubuntu, I am taken back to the GRUB menu! From the GRUB menu, I now select Ubuntu, which results in a fully functioning touchpad!

When using Xinput, I get nothing detected for Elantech but a notice

that I am using Xinput on Wayland. Under libinput, the Elantech mouse and touchpad are listed, but I have no mouse pointer or any response from the touchpad.

Colin Brammer

A There's a number of different models of Elantech touchpads, so it is difficult to make specific suggestions. It would appear that the touchpad is either being initialised or reset when you enter the firmware settings, so that is a good place to start.

A number of touchpads and keyboards use the i8042 driver module and you may need to pass settings to this. The laptop on which this is being typed has no functioning keyboard unless we pass a couple of options to the module when booting. You are in a more fortunate position in that you can at least type things to try to resolve the issue. Firstly, check that you are using the i8042 module with the following terminal command:

```
$ lsmod | grep i8042
```

If it returns any output, you can continue to experiment with the module. Try rebooting and adding these options to the kernels options: press **e** at the GRUB menu, go to the end of the line starting with 'linux', and add:

```
i8042.reset i8042.nomux i8042.nopnp i8042.noloop
```

If this helps, add the options to the **GRUB\_CMDLINE\_LINUX** in **/etc/default/grub** and run:

```
$ sudo update-grub
```

Once you get it working, you may want to experiment with the options to see which are actually needed to keep your touchpad working, and remove any unnecessary ones.

There are a couple of other things to try. Some report that the Synaptics touchpad driver interferes with this hardware, so you could remove any Synaptics packages (not *synaptic* – that is a useful package manager alternative). Finally, since entering the firmware menus helps, you could look for any firmware updates for your laptop from HP.



The details may vary, but all UEFI computers have an option in the firmware menus to turn Secure Boot on and off.

## Q Magick moments

I am trying to create a slideshow from a number of JPEG photos using *FFmpeg*. The issue I am facing is that the photos are of different sizes and I cannot get *FFmpeg* to work with all of them. There are a huge number of options for *FFmpeg*, which tells me two things: there must be one to do what I want, and I will never find it among the hundreds of other options. Can you suggest a solution?

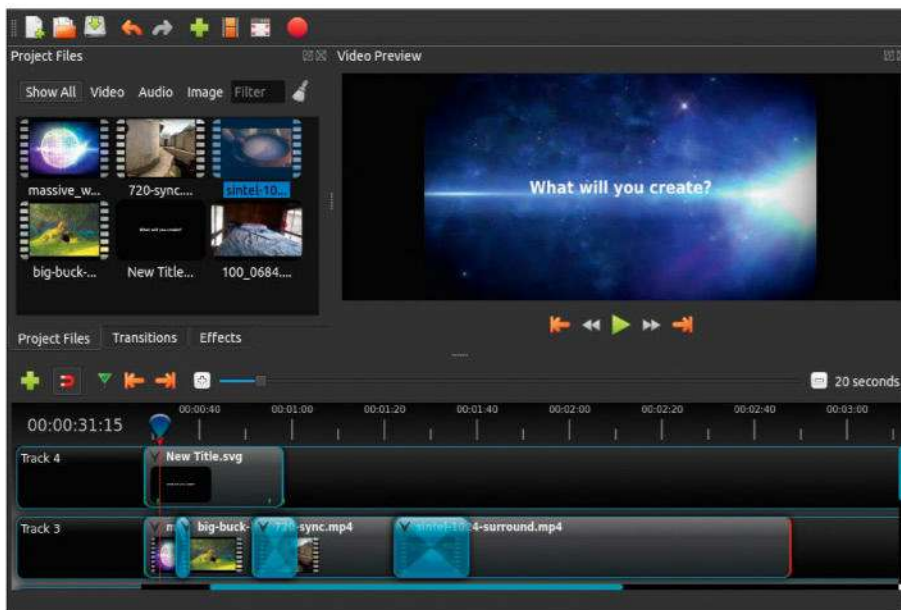
Peter Sanderson

A There are a few graphical programs to create slideshows from photos, such as *OpenShot* ([www.openshot.org](http://www.openshot.org)) or *Imagination* (<https://imagination.sourceforge.net>). As you're asking about *FFmpeg*, we assume you have a reason for doing it this way, possibly because there are too many photos to use a point and click option without getting RSI.

You do not say how you want the different sized images displayed. Do you want smaller images displayed smaller with a border or do you want all the photos to fill the screen? Also, do you want the final video's resolution to be the same as the largest photo, or do you need to resize everything to a lower resolution suitable for display on a monitor or TV? The latter seems more likely, so we will look at that.

Instead of trying to force *FFmpeg* to work with different sized images, it is easier to use a dedicated image conversion tool to create images of the same resolution first, then pass those to *FFmpeg*. The ideal tool for this is *convert*, part of the *ImageMagick* package. Let's say you have a text file called **photos.txt** listing all the photos you wish to add, one per line and in order. To resize them all to 1,920x1,080, suitable for a TV slideshow, you would use:

```
$ mkdir -p converted
$ rm -f ready_photos.txt
```



OpenShot can create slideshows, but it may not be so easy to use if you want to add gazillions of photos.

```
$ for PIC in $(cat photos.txt); do
  convert $PIC -resize 1920x1080 -gravity
  center -background black -extent
  1920x1080 converted/$PIC
  echo "file 'converted/$PIC'\nduration 5"
>>ready_photos.txt
done
```

If you are using *ImageMagick* 7.0 or later, you should use *magick convert* instead of *convert*, although the above command still works for now. It resizes the image to fit in a 1,920x1,080 space. This is normally narrower than 1,920 pixels as photos usually have a taller aspect ratio than TV screen. So, we use the **-extent** option to pad it out to the right size, with the image centred on a black background. Feel free to use another colour for the background or even an image file. The *convert* command creates the resized images, the following line puts the names of the converted photos into a file suitable for sending to *FFmpeg*. For each image there are two lines: the filename and the

slide duration. You can edit this prior to sending it to *FFmpeg*. You send it like this:

```
$ ffmpeg -f concat -i ready_photos.txt
-pix_fmt yuv420p -y video.mkv
```

The **-f concat** option tells *FFmpeg* that the input file is a list of files in a specific format used by its **concat** operator; the *FFmpeg* wiki has more on this. The **-pix\_fmt** option ensures compatibility with less than compliant players. Any other *FFmpeg* options can be added as you wish.

## Q Too many files

I have a huge media collection with hundreds of thousands of images and tens of thousands of video files. There are a lot of duplicates within all of these. The files are spread out over quite a few hard drives, both internal and external. I would like to be able to scan these files and determine which ones are duplicates, so I can then decide what to do with them.

I'd like the scan to be very thorough, not just a filename scan. I'd prefer

## » A QUICK REFERENCE TO... FILE LINKS

Linux makes extensive use of file links. This is a way to have the same file (or directory) appear in two or more places, or under different names. You could simply copy the file to each location, but this has a couple of disadvantages: it wastes space, and if the original file is updated, the copy is still the old version. There are two types of link: hard links and symbolic links,

aka symlinks. Technically, every file is a hard link, as this is a link from the file's name to the data structure on the disk that holds its contents. When referring to hard links, we tend to mean extra links, so a file has more than one name. Hard links have some limitations, most importantly that they only apply to files, not directories, and that all links must be on the same

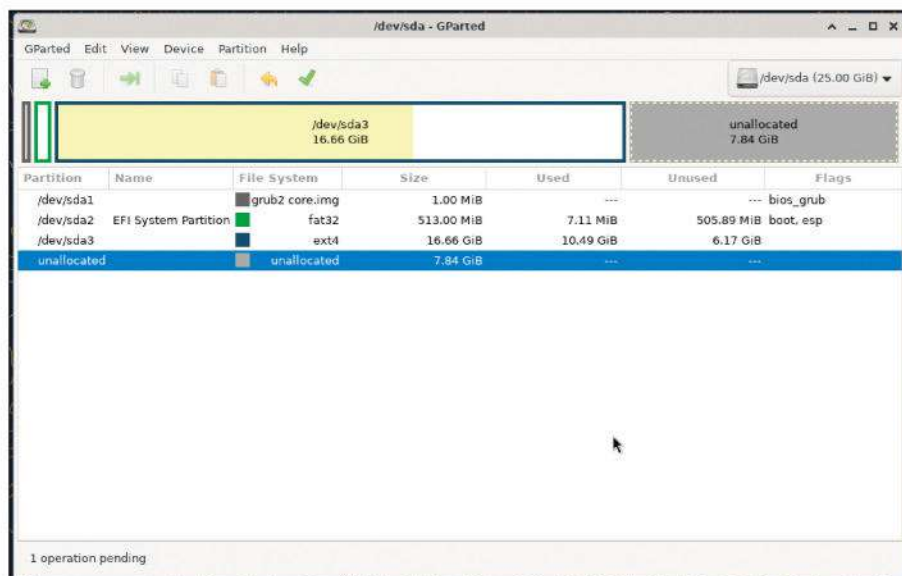
filesystem. However, you can delete any one of the files without affecting the others.

A symbolic link is far more flexible and you see them throughout a typical Linux system. You can see them when using the **-l** option with the *ls* command and they show up in most file managers. A classic example can be seen in the **/usr/src** directory, where **/usr/**

**src/linux** is a symlink to the source for the current kernel. So, you can have more than one kernel installed and anything that needs the sources can look in **/usr/src/linux** to find the current one.

Symlinks can be for directories or files, and can traverse filesystems. However, if you delete the 'real' file, the symlinks are left dangling, pointing to nothing.





GParted provides simple point-and-click partition management for desktop Linux systems, but using the command line is sometimes preferable.

multiple types of scans or verification to determine if these are the same file or not: filename, file size, dimensions of image or video file, or a digital fingerprint scan of the files.

**Zak Hooper**

**A** There are several tools for finding duplicates. One we use regularly is *Fdupes* (<https://github.com/adrianlopezroche/fdupes>). It works well with large collections as it eliminates potential duplicates based on the simplest tests first, such as file size. Only files that have the same size are compared more closely. Filenames are not considered; only the contents of the files matter.

All of the drives you need to scan should be mounted at the same time for this to work, then you can generate a list of duplicate files with:

```
$ fdupes -r /mnt/drive1 /mnt/drive2 ...
>dupes.txt
```

This generates a text file with all the duplicate files listed one per line, in separate groups. The **-r**, or **--recurse**, option scans all files on each given directory. You may also want to use the **-1**, or **--sameline**, option to output each set of duplicate files on a single line. If you are going to scan the same files more than once, add the **-c**, or **--cache**, option. This generates a cache of the digital signatures that *Fdupes* uses to compare files with the same size. Adding this option means that subsequent scans of the same directories will be faster. Should you need to remove it to start afresh, the cache is stored in **~/.cache/fdupes**.

If working with multiple directories, be careful not to specify the same directory twice. If you do that, *Fdupes* scans that

directory twice and lists all the files in it as duplicates of themselves!

By default, *Fdupes* only picks up true duplicates, that is two separate copies of the same file stored on your disks. If you also wish to find files that are duplicated by hard links, where there is only one copy of the data but more than one directory entry pointing to it, you should add the option **-H** or **--hardlinks**. Removing hard links can tidy things up but it will not free up any significant space..

### **Q Deleting without deleting**

If I delete a partition using *Fdisk*, will that delete the data on this partition? I want to use this in order to increase the partition size (after I have increased the filesystem size).

**Ben Ellis**

**A** The partition table contains information on the location and type of partitions on the disk. Deleting a partition from the table removes that information but the actual data on the partition is not touched unless you write to that area of the drive. What you describe is in fact the standard method of increasing the size of a partition, assuming you are leaving the start position unchanged and only moving the end. First you delete the old partition, then you create a new one starting at the same place but finishing further into the disk, then save the partition table. At this point, you may see a warning that the kernel is still using the old partition table. You can deal with this by running:

```
$ sudo partprobe
```

You probably already have *partprobe* on your system; if not, install the **parted**

package. Alternatively, you could simply reboot – but be careful about rebooting if you are resizing a system partition, especially your root partition.

Once you have resized the partition on the disk, the filesystem needs to be resized, using the appropriate tool. For ext2/3/4 filesystems, this is *resize2fs*, using it like this:

```
$ sudo resize2fs /dev/sdb3
```

If you do not give a size to *resize2fs*, it resizes the filesystem to fill its partition, which is what you want. Note that this is the opposite to the order you describe; you resize the partition first and then the filesystem.

To reduce the size of a filesystem and partition, you do the steps in the opposite order. First you reduce the size of the filesystem, then delete and recreate the smaller partition. For safety, it is advisable to reduce the filesystem to a smaller size than your intended partition size, then run *resize2fs* without a size as a final step to have it match the new partition.

If you are working with a desktop system, it may be simpler to install and use *GParted*. It performs the same steps, but gives you a graphical interface where you can simply click and resize partitions, and let *GParted* do the donkey work.

Resizing partitions like this, particularly enlarging them, is safe, but nothing is risk-free. A random hardware crash or power failure at the wrong time could cause data loss, or at least create a lot of work in data recovery. It is always a good idea to have recent and verified backups before doing anything like this. **LXF**

### **GET HELP NOW!**

We'd love to try to answer any questions you send to **answers@linuxformat.com**, no matter what the level. We've all been stuck before, so don't be shy. However, we're only human (although many suspect Jonni is just a dream), so it's important that you include as much information as you can. If something works on one distro but not another, tell us. If you get an error message, please tell us the exact message and precisely what you did to invoke it.

If you have, or suspect, a hardware problem, let us know about the hardware. Consider installing *hardinfo* or *lshw*. These programs list the hardware on your machine, so send us their output. If you're unwilling, or unable, to install these, run the following commands in a root terminal and send us the **system.txt** file, too:

```
uname -a > system.txt
lspci >> system.txt
lspci -vv >> system.txt
```

# Mailserver

## WRITE TO US

Do you have a burning Linux-related issue that you want to discuss? Write to us at *Linux Format*, Future Publishing, Quay House, The Ambury, Bath, BA1 1UA or email [letters@linuxformat.com](mailto:letters@linuxformat.com).

## You're no editor

As a new Windows escapee, I'm finding the switch to be a little confusing, though easier than I expected. I was digging around my Debian desktop and couldn't seem to find anything like the Windows Registry editor. I came across the *dconf* editor – is that the same thing? It seems to offer the same features.

**Rodger Hemmingway**

Neil says...

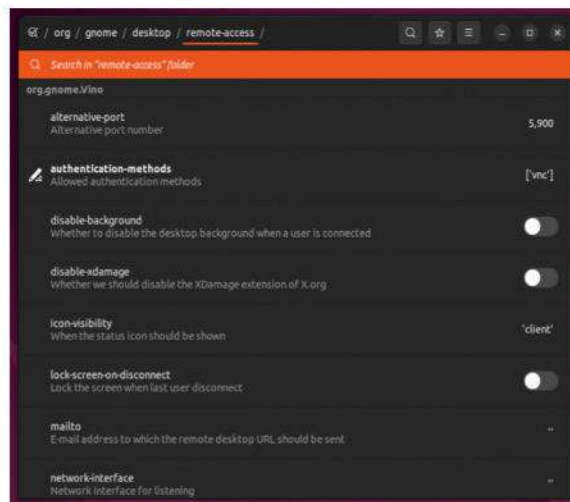
The advantage Microsoft has with Windows is that it's the sole arbitrator over development, so it can say, "Hey, this is the central store for all configuration data," and – boom – it happens. So, a single Registry editor is possible for Windows. The open source Linux ecosystem is far more tricky. The *dconf* editor is, in fact, the Gnome desktop's configuration editor. It lets you tweak all the various Gnome tools, but it's only relevant to the Gnome desktop. All the other system settings – such as to Cinnamon if you switch desktops – are controlled differently. You could theoretically create a project to handle all of these but it would be a monumental effort to implement and maintain, and those efforts are better used elsewhere!

## Think about it

I have come up with an idea for NPUs. They should be made three-dimensional, like the human brain.

**Ian Learmouth**

Not that there's much we can do with them yet, but here's a bunch of NPUs to look at.



You can edit low-level configuration data with *dconf* but it's only applicable to the Gnome desktop and apps.

Neil says...


NPUs are just glorified GPUs – all the 'thinking' is done in the software language model, which is statistical abstractions of huge amounts of processed data. So, a 3D NPU wouldn't really function any differently or better. Not that LLMs actually think; they're just outputting the most likely next character/pixel, which is why they make basic mathematical errors and hallucinate "facts".

## Moanbuntu


This is definitely a moan, but I'll try to be positive at the end. I'm running Ubuntu 16.04 very happily. I thought, however, that I had better get up to date, so when I saw the mag with Mint on the cover, I decided it was time to get up to date. I bought the issue (LXF319), and discovered that Mint was based on Ubuntu, so I back-ordered the mag (LXF315) that covered it.

Then issues started. Firstly, the mag did not make it clear that the change to the BIOS was permanent (I read it as temp, for the non-install running). Fortunately, I just about know what I'm doing, so that was irritating, rather than anything else. Messed about, as required, and got Ubuntu


CREDIT: Intel, Balena



## Helpdex







installed. Couldn't then load my bookmarks into *Firefox*, and couldn't get any email client to accept my emails (I run POP3, so they're stored locally).

On to Mint, or so I thought. Mint doesn't download as a straight ISO, so the startup disk creator for 16 doesn't know what to do (it coped with Ubuntu 24!). More annoyance, compounded by the impossibility of getting *Etcher* to run. The download was fine, but no executable that would run, and no help on its site. One day later, I've got nowhere, and got a flash drive I don't want. Can you let the community know that they really need to look at this whole install area much more thoroughly if they want to move users to Un\*x from anywhere else. Spinning a disk (16 and I think 18) is fine, but all this is ridiculous. Nobody is going to take the time. You might also look at migration paths, as well as clean installs, to see what issues that throws up. As for me, I'll stay on 16.04. It should see me out.

**Hugh Stirling**

Neil says...

I wouldn't say you're moaning – if you're encountering genuine usability issues, it's good to know about them, as it helps explain things better in the future. *Etcher* is a fine example; the ZIP content has no explanation as to what you need to do with it – standard advice is to get it via the app store. I would point out, though, that Mint is an ISO file by default, so that's intriguing.

What version of *Firefox* are you running? There was a format change years ago but it hasn't been a problem for over a decade. It's not something I've dealt with, but you can have different profiles, and data can be retained in them. Here's a link: <https://mzl.la/3BJEHOP>

I'm also curious what email client you were using. *Thunderbird* has been the go-to client for years and that supports local files and POP3 access. I'm guessing the project you were using was abandoned much like Ubuntu 16.04 that you're still using... **LXF**

## » LETTER OF THE MONTH

### Firefox

A recent *Firefox* update has, once more, disrupted my browsing. I have used *Firefox* for many years, appreciating the Private Browsing mode that has, so far, protected me from the annoying adverts and pop-ups that plague others. However, my install of Ubuntu 20.04 updates *Firefox* without warning and crashes all the Private Browsing tabs. This is just as bad as a certain popular OS from Redmond. What is wrong with the 'download the new and switch at reboot' approach of other software? Is this the doing of Mozilla, or is it the drafted Snap approach to app deployment?

**John Gillham**

Neil says...

This does sound like it's a Snap issue. Although Mozilla is packaging the Snaps these days, I believe, you'd think that would cut out some of the issues. There was a recent security update to fix an in-the-wild exploit, so sometimes updates can't really wait, but I am with you on the point that these should really always be user-triggered.

We've read that trying the following can reset Snap issues:

```
$ sudo killall snap-store
$ sudo snap refresh
```

As for Snap updates, I believe they're auto-updated every day. You can implement delays and controls individually or system-wide (I wouldn't recommend that). This link has details: <https://snapcraft.io/docs/managing-updates#heading--controlling-updates>

Perhaps try this:


```
$ snap refresh
--hold=<duration>
<snap1> <snap2>
$ snap refresh
--hold=24h firefox
```

Possibly a better option from that link would be the following – it asks the system to only refresh snaps between 4.00am-7.00am, and 7.00pm - 10.10pm:

```
$ sudo snap set
system refresh.
timer=4:00-
7:00,19:00-22:10
```



Naughty Snaps! Stop breaking the nice reader's Firefox tabs.



shane\_collinge@yahoo.com



SUBSCRIBE Save money today!

# SUBSCRIBE TO LINUX FORMAT MAGAZINE SAVE 50%\*

[www.magazinesdirect.com/LIN/H23A](http://www.magazinesdirect.com/LIN/H23A)

Call 0330 333 1113  
and quote H23A

**NEW!**

Digital access to  
130+ issues when  
you subscribe  
to print!\*\*





Save money today! **SUBSCRIBE**



**OUTSIDE THE UK?**  
Turn to page 67 for more great subscriber deals!

» **NEW:** Improved digital back-issue access!\*\*



Get **instant access** on your browser, or Android and Apple iOS devices back to issue 181 (March 2014) with 1,000s of tutorials, interviews, features and reviews.

## » CHOOSE YOUR PACKAGE!

Keep your subscription rolling!

Subscribe today and save 50%!\* + All-new digital access\*\*

### PRINT EDITION



Only  
**£21.09**

for 6 months  
6 months of Linux Format delivered for £21.09!

### DIGITAL EDITION



Only  
**£22.71**

for 6 months  
6 months of digital Linux Format for £22.71!

**SAVE!**  
**50%\***

Terms and conditions: **Offer closes 7th January 2025.** Please allow up to 6 weeks for the delivery of your first subscription issue (up to 8 weeks overseas). The full subscription rate includes postage and packaging. \*Savings are based on the cover price. Payment is non-refundable after the 14-day cancellation period. \*\*Access to the digital library will end with your subscription on print. For a digital subscription, once terminated you only keep the issues you have paid for up to that point. For full terms and conditions, visit [www.magazinesdirect.com/terms](http://www.magazinesdirect.com/terms). For enquiries and overseas rates please call: +44 (0) 330 333 1113. Lines are open Monday-Friday 8:30am-7pm, Saturday 10am-3pm UK time (excluding bank holidays) or email: [help@magazinesdirect.com](mailto:help@magazinesdirect.com). Calls to 0330 numbers will be charged at no more than a national landline call, and may be included in your phone provider's call bundle.

# YOUR DIGITAL ISSUE ACCESS

Linux Format print subscribers can now access digital back issues two ways!\* Who's a lucky bunch of readers?!



» As a big thank you for subscribing to *Linux Format Print Edition*, we've always offered digital access to past issues. It seems The Management thought it was a good idea and has rolled out a new system using Pocketmags!

This means there are now two ways of accessing LXF back issues for subscribers. The new Pocketmags service goes back to March 2014, including over 140 full issues, while our older PDF archive goes back to LXF66 (a crazy 8-bit filling 254 issues), though issues before LXF120 have had covers and ads stripped. We've outlined how to access both below. Enjoy!

\* Terms and conditions apply – see page 17 for full details.



## FANCY NEW POCKETMAGS

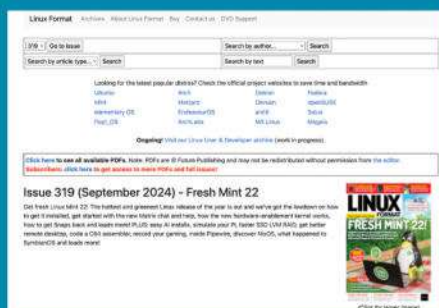


**1 Activate your access**  
Head to [www.magazinesdirect.com/digital-instructions/](http://www.magazinesdirect.com/digital-instructions/) and click the Start Reading link that'll transfer you to the Pocketmags activation page.

**2 Create an account**  
Enter the email that was used to create the subscription. You'll be prompted to create a Pocketmags account so you can access the library.

**3 Activate your sub**  
On the Activate A Subscription page, select Linux Format from the pull-down menu and click on Go To Activation Page. Select Library – happy browsing!

## THE CLASSIC LXF PDF ARCHIVE



**1 Head to Linux Format**  
You heard the man – browse to [www.linuxformat.com/archives](http://www.linuxformat.com/archives) using your favourite web browser. Where it says Subscribers: Click Here, do just that.

**2 Use your subs number**  
Log in using your subscriber number and surname. You can find the number above your name and address on the address label.

**3 Access your issues**  
Use the issue drop-down menu to access past and current issues. You can also search by author and keyword or article types.

**DON'T MISS!**  
Includes 5 years of Linux User & Developer issues

Visit [www.magazinesdirect.com/linux-format](http://www.magazinesdirect.com/linux-format)



# TP-Link BE3600

Is it still Wi-Fi 7, asks **Brandon Hill**, if there's no 6GHz?

## SPECS

**Wi-Fi:** Wi-Fi 7

**Bands:** 2.4GHz:

2x2 (Tx/Rx) up

to 688Mb/s

5GHz: 2x2

(Tx/Rx) up to

2,882Mb/s

**Ports:** 1x

2.5Gb/s for

WAN, 1x

2.5Gb/s for

LAN, 3x 1Gb/s

for LAN,

USB 3.0

**USB NAS:**

NTFS, exFAT,

HFS+, FAT32,

FTP, Samba,

DLNA

**Parental:**

Content filter,

family time,

bedtime, off

time, time

rewards, time

ranking, internet

pause, reports



**W**i-Fi 7 client devices pale in comparison to their Wi-Fi 6E and Wi-Fi 6 counterparts, and Wi-Fi 7 routers can cost hundreds of dollars, with high-end mesh systems easily surpassing £500. Some router manufacturers are combating the pricing issue by taking the unorthodox approach of excluding one of the pillars of the Wi-Fi 7 standard, such as the 6GHz band, to appeal to more value-conscious buyers.

The TP-Link Archer BE3600 has four antennas, clustered around the rear corners of the router. The business end of the BE3600 features one 2.5Gb/s port for WAN and one 2.5Gb/s port for LAN. There are three 1Gb/s LAN ports to round out the wired networking portfolio. There's also a USB 3.0 port for connecting a printer or storage device. This is a 2x2 Wi-Fi 7 router, which limits its wireless performance to a maximum bandwidth of 3,570Mb/s (688Mb/s + 2,882Mb/s).

## Nice and breezy

The BE3600 is a breeze to set up. After powering on, point a browser to <http://tplinkwifi.net>, which starts the configuration process. Choose a new admin password and options to configure the Wi-Fi radios. The software interface for the Archer BE3600 is similar to what we saw with TP-Link's more powerful Archer GE800 tri-band Wi-Fi 7 router.

There are five primary tabs at the top of the homepage: Network Map, Internet, Wireless, HomeShield and Advanced. It's lacking the Game Center tab from the GE800 for obvious reasons. One other change of note is that, by default, the BE3600 has a lighter, more cheerful turquoise and white colour scheme for the user interface.

*iPerf3* tests on the 2.4GHz band at six feet were strong – the BE3600 managed 214Mb/s with no traffic and 205Mb/s when congested. Both figures put the router well ahead of the competition, even the pricier Archer GE800 and Asus RT-BE96U. However,

performance at 25 feet took a dramatic turn, falling to just 51Mb/s and 45Mb/s, respectively, putting it in last place. It was a similar story with the 5GHz tests, with the BE3600 putting up a respectable 1,121Mb/s at six feet with no traffic and 1,030Mb/s when congested. However, *iPerf3* performance fell to just 382Mb/s (uncongested) and 374Mb/s (congested) at 25 feet.

At times, it can be difficult to wrap your head around the premise of the TP-Link Archer BE3600. It's labelled as a Wi-Fi 7 router, but it doesn't support the full Wi-Fi 7 spec – namely, it's missing the 6GHz band, and performance suffers as a result. In addition, since this is a 2x2 Wi-Fi router, it offers half the theoretical performance of other dual-band Wi-Fi 7 peers, such as the Asus RT-BE88U.

If you're not big into gaming and just want a no-frills router to support a 1Gb/s or lower internet connection, there's nothing wrong with the BE3600. Just realise that you're leaving a lot of performance on the table compared to fully featured Wi-Fi 7 routers. If you want added performance without breaking the bank, consider looking at budget Wi-Fi 6E routers. **LXF**

## VERDICT

**DEVELOPER:** TP-Link

**WEB:** [www.tp-link.com/uk](http://www.tp-link.com/uk)

**PRICE:** £110

**FEATURES** **6/10**

**PERFORMANCE** **5/10**

**EASE OF USE** **9/10**

**VALUE** **7/10**

Cheap, but you also don't get the full spec of features or performance of 'true' Wi-Fi 7 routers, so what's the point? Though 2.5Gb/s ports do help.

» **Rating 6/10**

# Ubuntu 24.10

**Nate Drake** brims with excitement at Canonical's latest offering. Will this distro be a plumed Oriole or just plain 'Orrible?

## IN BRIEF

An oriole is a brightly-coloured bird with a rich, whistling song. This carries over into Ubuntu 24.10, which has an upgraded installer, a new Gnome desktop and an experimental *Security Center*.

## SPECS

**CPU:** 2GHz dual-core  
**Mem:** 4GB  
**HDD:** 25GB  
**Builds:** x86-64, ARM640, RISC-V, ppc64le (POWER8 and later), s390, ARMHf (ARMv7 + VFPv3-D16)

**T**wenty years ago, in October 2004, the very first stable release of Ubuntu (4.10) was made available to the public.

It represented the fruits of the labour of Mark Shuttleworth and a small team of Debian developers. Since that time, with one exception, new releases of Ubuntu have been issued every six months. Version numbers correspond to the current year and month of release.

This has culminated in the latest version (24.10). In true Ubuntu tradition, it bears the alliterative code name of an animal (this time, it's Oracular Oriole).

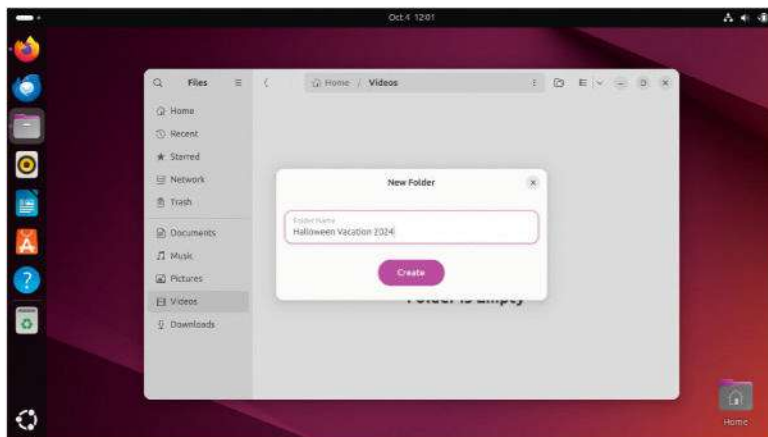
As one of the most popular Linux distros, Ubuntu supports a wide range of architectures. The AMD64 ISO is a rather weighty 6.7GB. This is around a 1GB increase over Ubuntu 24.10, which was roughly the same size as a Windows 11 ISO.

If you decide to test Ubuntu 24.10 in *VirtualBox*, as we did, you may see an error message on startup: `vmwgfx seems to be running on an unsupported hypervisor`. In our tests, we found this disappeared after installing the Guest Additions.

From going through the *Flutter*-based installer, there seems to be no apparent changes from the previous version of Ubuntu. Users can select their language, as well as configure accessibility options and keyboard.

The installer also still checks whether it's up to date, and if not offers the choice to upgrade before continuing with the setup. You need to restart the installation if you choose to do this.

After opting for an automated installation, we noted that Ubuntu's setup also still offers a choice between the Default selection of apps, to include essentials such as



Oracular ships with the latest Gnome 47, with support for accent colours, a customisable file manager and overhauled system dialogs.

the web browser, or the Extended selection to include additional programs, such as office tools.

The Extended option no doubt accounts for the large size of the ISO. We can't help but feel it might be more sensible to reduce the distro to the Default apps but offer to download and install others during setup if users opt for the Extended version.

After opting for the Extended installation and creating a user account, we found that setup completed in just over 11 minutes. Upon restart, we noted that in commemoration of Ubuntu's 20th anniversary, the OS now plays the original startup sound. (This can be disabled via *System Settings*.)

The system also displayed the current release notes, which state that the system also has a default 'warty' brown accent colour.

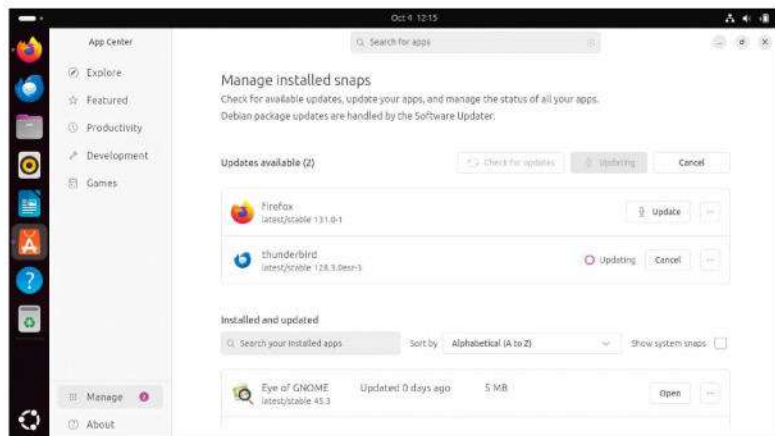
This is significant, because Ubuntu now ships with the latest Gnome 47. Users of this desktop environment have been excitedly talking about how it now supports accent colours.

Of course, this is old hat for Ubuntu users, given the OS has supported this feature since 2022. Still, this should ensure better compatibility across third-party applications. Accent colours can be accessed in the Appearance section of *System Settings*. We used it to add a purple tinge to our Ubuntu desktop.

Gnome 47 has also introduced improved system dialogs that make use of floating buttons. They have also been overhauled to function better on narrower displays.

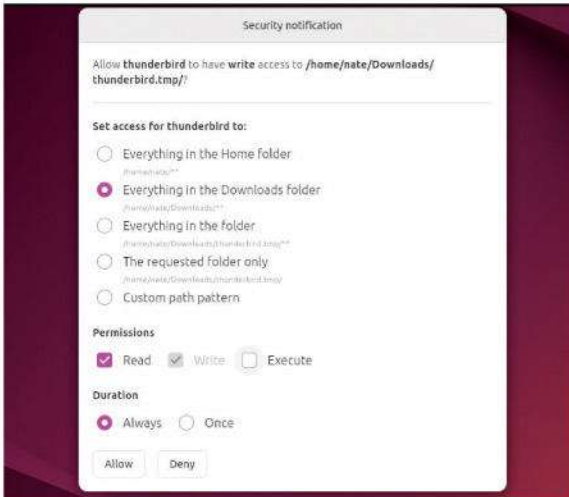
Other visual changes are more subtle. For instance, when you now right-click an icon in the dock, the context menu lists the application name in a header label. If the program in question is a Snap package, you can also now choose App Details to view more information in the *App Center*.

The *App Center* itself has also received some minor tweaks. Upon launch we noticed that featured Snap packages are now displayed more prominently. There's



The Manage section of App Center now supports monitoring updates (progress is shown in the dock) and removing installed Snap packages.





The experimental Security Center requires you to set permissions for Snap apps. You can allow these once only or permanently.

also a new Art & Design category, which includes applications such as *GIMP* and *Blender*.

The *App Center*'s Manage section now includes information on installs in progress, and displays alerts when Snaps are running. Users can now also easily uninstall Snaps from here.

When we used the Manage section to update installed Snaps, we noticed a new quirk of the Ubuntu dock: it now displays progress bars underneath the icon of the app that's currently updating. This is a very small addition but reflects the care Ubuntu developers put into making the interface as intuitive as possible.

The *App Center* also now supports third-party DEB installation. This is a long-awaited feature, so we tried it out by downloading the DEB for *Proton VPN*. By default, DEB packages still seemed to open in *File Roller*. Still, when we right-clicked and chose Open With App Center, we were able to install the package without issue.

The *Nautilus* file manager has also been overhauled thanks to the GNOME upgrade. Crucially, users can now remove bookmarks to common folders like **Documents**. The Trash has now been placed at the top section.

The GNOME release notes state that all internal drives are now listed on the sidebar, as opposed to a shortcut to Other Locations, but this wasn't in our version, which we put down to the fact that it is a development release.

GNOME 47 has also slightly tweaked the *Disk Usage Analyzer* interface to make it visually cleaner. We fired this up to discover that our Extended install took up almost exactly 12GB on the virtual hard disk.

Another major upgrade is the inclusion of a new *Security Center* app. At the time of writing, there's just one experimental feature: Require Apps To Ask For System Permission. After enabling this, each time we launched a Snap app for the first time, Ubuntu served a security notification via its new prompting client. This enables you to set strict permissions to allow read or write access for specific folders. You can also set these permissions to apply only once or always. You can also choose More Options to specify a custom path, as well as assign execute permissions.

After granting *Thunderbird* access to the entire **Downloads** folder in our test machine, it launched immediately. However, the prompting client didn't appear any of the times we launched *Firefox*.

## » 20 YEARS OF UBUNTU

The Ubuntu developers who released Warty Warthog in 2004 may have found it difficult to imagine that they'd be breaking out the 20th birthday candles for Ubuntu in October 2024.

The name derives from a Bantu word, which can loosely be translated as "community spirit". This reflects Canonical's collaborative approach with the community to OS development.

For example, in 2005, the developers began the Shiplt programme, which supplied free Ubuntu CDs to anyone who requested them. It's hardly surprising that in the same year the OS won multiple awards for being the favourite distro for Linux users.

In 2023, Canonical proved Ubuntu isn't going anywhere any time soon by introducing Ubuntu Pro, which, combined with LTS releases, can extend OS support for up to 12 years. Here's to many more!

Command line lovers will also be pleased to discover that *Oracular* now ships with the latest version of *APT* (3.0). This is a huge improvement over the traditional wall of text that appeared in the previous version. Dependencies now appear in neat, coloured columns. *APT* also now displays suggested packages and a quick summary before prompting users to proceed with install.

From reading the release notes, we saw that devices using Nvidia GPUs default to Wayland sessions for the first time. Upon signing out, we noted that you can still select Ubuntu On Xorg from the login screen.

We ran *System Monitor* to check how well Ubuntu performs relative to the stated system specs. With only the monitor running, the system consumed around 2.5GB of RAM and 1% of our 2GHZ virtual CPU. **LXF**

```
nate@ubuntu2410: ~$ sudo apt install gimp
Installing:
gimp

Installing dependencies:
gimp-data libcdt5 libgts-bin libmypaint-common
graphviz libcglyph6 libgvc6 libopenexr-3.1.30
libamd3 libcholmod5 libgvpr2 libpathplan4
libann0 libgegl-0.4-0t64 libmath-3.1-29t64 libumfpack6
libbabi-0.1-0 libgegl-common liblab-gamut1
libcamd3 libgimp2.0t64 libimg2
libccolamd3 libgts-0.7-5t64 libmypaint-1.5-1

Suggested packages:
gimp-help-en | gimp-help gimp-data-extras gsf fonts graphviz-doc

Summary:
Upgrading: 0, Installing: 26, Removing: 0, Not Upgrading: 0
Download size: 29.4 MB
Space needed: 147 MB / 88.3 GB available

Continue? [Y/n]
```

## VERDICT

**DEVELOPER:** Canonical Ltd

**WEB:** <https://ubuntu.com>

**LICENCE:** Mainly GPL

FEATURES	10/10	EASE OF USE	10/10
PERFORMANCE	10/10	DOCUMENTATION	10/10

Like Tommy the Pinball Wizard, Canonical never seems to fall. The latest release is powerful, easy to set up and speedy.

» **Rating 10/10**

Ubuntu now comes with *APT* 3.0, which neatly displays data like package dependencies when you try to install software.

# Kali Linux 2024.3

**Nate Drake** explores the latest version of the distro beloved by both hackers and Mr Robot fans the world over.

## IN BRIEF

The team at Offensive Security has done it again, with better chip support, an expanded arsenal of tools and an upgraded kernel to keep even the most picky pen testers happy.

## SPECS

**CPU:** 1GHz  
**Mem:** 2GB  
**HDD:** 20GB  
**Builds:** x86, x86-64, armel, armhf

When reviewing the Devuan-based distro Expirion this month (*opposite page*), we noted the developer had created a spin of the OS based on Kali, minus the usual tools for “those who’ve watched too much *Mr Robot*”.

Although there’s no such a thing as too much *Mr Robot*, Kali does feature prominently in the series. This is because it isn’t intended as a daily driver but as a convenient way for hackers to make use of multiple pen-testing tools.

Kali is based on the latest testing version of Debian (Trixie). According to the release announcement, Kali now supports the Qualcomm Snapdragon SDM845 SoC.

Kali’s online support is so comprehensive, it even puts Canonical to shame. The documentation is well written, easy to navigate, and supplemented by official forums and a Discord channel.

The recommended installer ISO weighs in at 4GB. However, there is a 485MB netinstaller, as well as a live version and a 13GB ISO of everything for complete offline setup. We initially chose the Standard installer. The ISO boots to a text prompt asking you to choose your language, time zone and keyboard layout.

The installation media is scanned, at which point you are greeted with more text prompts to set a hostname, domain name, username and password.


The partitioner then starts, offering various methods, such as using the entire drive. System encryption is supported, and you can also set separate partitions for `/var`, `/home` and `/tmp`.

The text installer offers a variety of desktop environments, including Xfce, Gnome and KDE. After choosing the default options (Xfce plus recommended tools) and the location for the GRUB bootloader, our install completed in just under 10 minutes.

Overall, Kali setup doesn’t hold your hand in the same way as a graphical installer would. But its target market is users who are very familiar with command line tools. The text interface is also simple to navigate.

On restart, the desktop loaded in seconds. We launched the task manager to check resource usage against the requirements listed in Kali’s documentation. At rest, the system consumed around 785MB of RAM and an average of 2% of our 1GHz virtual CPU. Kali’s support pages caution that some intensive applications can consume far more system resources – for instance, the *Burp Suite* needs around 8GB.

The release notes mention 11 new tools in Kali, such as the excellent note-taking tool *Obsidian*. However,



Kali contains a default selection of popular hacking tools. You can add more by installing metapackages via the command line.

when we searched for these via the *Application Finder*, they weren’t installed.

From checking the excellent online documentation, it seemed that during setup, while we’d selected the Default and Top 10 toolset, the latest ones weren’t included. As Kali has no official package manager, we confirmed this by using the bundled *kali-tweaks*, which can manage the OS’s so-called metapackages.

Upon using the terminal to try to install a new metapackage and one of the tools, in both cases we were told they didn’t exist. The same thing happened when we tried to install *Synaptic*.

Undeterred, we opted to try out the 4.3GB live ISO instead. This spares users going through setup, although it does support USB/encrypted USB persistence. It also deploys Xfce.

The new tools also weren’t present here, but we were able to run updates and install *Obsidian* successfully via the terminal. The live version of *kali-tweaks* also offered a greatly extended selection of metapackages, including *kali-linux-large*, which we used to install the extended default tool selection. **LXF**

## VERDICT

**DEVELOPER:** Offensive Security

**WEB:** [www.kali.org](http://www.kali.org)

**LICENCE:** GPL 3.0

<b>FEATURES</b>	<b>10/10</b>	<b>EASE OF USE</b>	<b>9/10</b>
<b>PERFORMANCE</b>	<b>9/10</b>	<b>DOCUMENTATION</b>	<b>10/10</b>

Offers effortless setup and loads of well-categorised pen-testing tools. Installing extra packages could be a little easier.

» **Rating 9/10**



# Expirion Linux 5.0.2

**Nate Drake** casts aside systemd in favour of this lightweight Devuan-based distro. Is SysVInit the right way to go?

## IN BRIEF

We always admire distros maintained by a single developer; this fork of Devuan offers superior performance and carefully curated software. Unfortunately, there are bugs both in setup and system settings.

## SPECS

**CPU:** 1GHz  
**Mem:** 1GB  
**HDD:** 8GB  
**Builds:** x86\_64

As most Linux users know, the introduction of *systemd* divided the community. On one hand it can provide faster boot times and better service management, not to mention consistency across distros. Nevertheless when it was integrated into Debian 8, a group of developers felt *systemd* went against the Unix philosophy of “doing one thing well”. In 2014, they announced the creation of their own fork, Devuan. This is *systemd*-less, instead relying on components such as *init*, *eudev* and *elogind*.

This latest version of Expirion is itself a fork of Devuan 5 (Daedalus) SysVInit. The developer admits he kept the code as close to the original Devuan as possible, as well as retaining the default desktop environment, Xfce.

There's also a spin that ships with Gnome, as well as one based on the current testing branch (6.0) of Devuan. Another spin is based on Kali minus any pen-testing tools. On SourceForge, the creator admits he can't see any real use for this, but the main benefit is for people who've been watching “too much *Mr Robot*”.

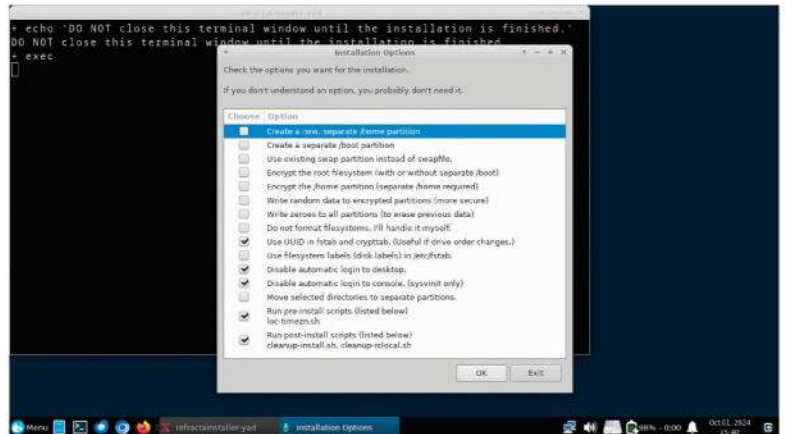
After downloading the 2GB Xfce spin of version 5.02, the live desktop booted in under 12 seconds. Here we noted that Expirion uses the installer for another Devuan-based distro, Refracta.

It's less intuitive than options such as *Calamares*. There's a separate terminal window to the main GUI. You also need to choose setup options, like encrypting the root partition, by ticking various boxes. We decided to proceed with the default options checked,

We were then prompted to open *Gdisk* or *Gparted* to prepare the partitions. After skipping this, we were able to select the correct partition and filesystem (ext4). The installer then asks you to specify your region and keyboard layout. Once done, files copied over in around five minutes. Finally, the installer prompts you to set a username and password, and whether to allow the account *sudo* privileges.

Upon restart, Expirion booted to the GRUB command line. Despite our best efforts to load the kernel and boot from the installed partition, Expirion failed to load. Sadly, there's little documentation to walk users through the install process beyond a brief **Readme** file on SourceForge. The developer does, however, encourage users to contact him via his user account on the **Linux.org** forums.

Instead we booted the live ISO once again and fired up *System Monitor* to check how economical Expirion is with system resources. As we expected for an OS



The Refracta installer uses tick boxes to choose setup options. During our tests, however, it failed to write a /boot partition.

deploying Xfce, the RAM requirements were very forgiving: only 850MB were used at rest. Usage of our 1GHz CPU also hovered at only around 2%.

Another strong point in Expirion's favour is the sensible selection of bundled apps. Both the *Chromium* and *Firefox* browsers come preinstalled. Email is managed via *Thunderbird*. The OS also comes with the latest *LibreOffice* 24.8.2. We were also delighted to discover a healthy selection of media apps, including *VLC*, *HandBrake* and *Audacity*.

Further software can be installed via *Synaptic*, which we used to obtain the classic 2D point-and-click adventure *Beneath a Steel Sky*. The first time we did this, the packages failed to download, but it installed on the second try. Both the game and the corresponding *ScummV* also immediately appeared in the Games menu. We also had no trouble installing the command-line utility *cowsay* via the Xfce terminal.

The OS doesn't seem to ship with any custom wallpapers. We also noted that the Appearance section of *System Settings* froze when we tried to change the OS theme. **LXF**

## VERDICT

**DEVELOPER:** Charlie Locke **WEB:** <https://sourceforge.net/projects/expirion-linux/>  
**LICENCE:** GPL & BSD

<b>FEATURES</b>	<b>7/10</b>	<b>EASE OF USE</b>	<b>5/10</b>
<b>PERFORMANCE</b>	<b>8/10</b>	<b>DOCUMENTATION</b>	<b>2/10</b>

Expirion is blazingly fast with a good selection of apps. It's let down by a tricky setup and poor documentation.

» **Rating 6/10**

# Core Keeper

Keep the staff in the dark? The Management are ecstatic at the latest innovation to come from the crafting stone of **Christopher Livingston**.

## SPECS

Minimum

OS: Ubuntu

20.04 64-bit

CPU: Intel Core

i5 2300, AMD

Ryzen 3 1200

Mem: 8GB

GPU: GeForce

GTX 460,

Radeon HD

5850

Recommended

OS: Ubuntu

22.04 64-bit

CPU: Intel Core

i5 8400, AMD

Ryzen 7 2700X

Mem: 8GB

GPU: GeForce

GTX 1050 Ti,

Radeon R9

280X


**W**e wish we'd kept track of how often we exclaimed "What the heck is that?"

while playing *Core Keeper*. It must have happened a dozen times in the first dozen hours alone, such as when coming across a conch shell the size of a school bus, or stumbling upon a subterranean castle 20 times the size of our base. And let's not forget when a caterpillar as big as a passenger train came stampeding through the tunnel that was being explored. What the heck?

*Core Keeper's* got a lot in common with survival games such as *Valheim*, plus plenty of *Minecraft* DNA, and even a smidge of *Stardew Valley*. Bizarre sights, unexpected encounters, and delightful (though dangerous) surprises are just a few things we want from a procedurally generated co-op survival crafting sandbox, and those aren't the only important items *Core Keeper* checks off the list.

Colourful and cute as its pixellated world is, it gets more daunting and scary the deeper you go – yet despite all its dangers and some seriously powerful bosses, it also makes for a cosy and comfy base-building experience. Plus, it's got fishing. No wonder it was easy to sink more than 60 hours into it.

In *Core Keeper*, you (and your friends, since it supports up to eight players) find yourself stranded underground in a dimly lit cave filled with ancient and defunct technology. Awakening the slumbering core means finding a way to power it up, so pick a direction and just start digging. Chip away at dirt walls, harvest the materials you find, craft new weapons and gear, build a base, and dig far-reaching tunnels outwards from the



Core Keeper's lighting, effects and animations are genuinely beautiful.

core to other biomes, which include clay caverns, arid deserts and even a vast blue ocean.

## Digging in

Progress is familiar if you've played some survivalcraft before: mine early resources such as copper and smelt the ore into ingots to make better tools, armour and weapons, then dig deeper into the world to find sturdier stuff such as tin, gold and iron. But it's not all about whacking walls with a pick; one of the joys of *Core Keeper* is setting up resource production lines so you can get machinery to do the grunt work for you.

Finding a chamber with copper and iron boulders too hardy to be carved, we marked it on the map and came back once enough progress had been made to drilling technology. Powered by an electrical generator, the drills worked around the clock to dig out those nodes, while robotic arms grabbed the ore nuggets and dropped them on conveyor belts, which transported them through the tunnels we'd dug to our base.

There, another set of arms collected the ore and dropped it into a smelter to become ingots, where they were again picked up and placed in a chest next to our workbench. It took ages to set all this up, not to mention all the resources that had to be collected by hand to reach that level of automation – but it was super-satisfying to put into action. A world that had once been an intimidating network of darkened tunnels concealing the unknown eventually became an elaborate and efficient mining operation.

Combat isn't a particular high point of *Core Keeper*, and it begins as pretty crude hack-and-slash with the early threats you face, from bobbling little blobs of slime, to goblin-like Caveling tribes, to angry mushroom men. As you progress, you'll find and craft new weapons like a blunderbuss –



A boss fight goes well for the group of attackers.





which seems unfair to the mushroom men – enchanted swords, magic wands and, our favourite, the Grubzooka, which is like a rocket launcher made from a giant maggot. It's so powerful you don't need to dig through dirt and stone walls, just 'zooka them. There's also a tendency to blow yourself up, but it's mostly worth it.

There's also plenty of enjoyable details in the enemies you face, especially the Cavelings. At first they just seemed like mindlessly hostile goblins, but sometimes you come across one chipping away at a wall, mining for ores just like you do. We've even seen several of them sleeping, because even monsters need a rest. Spend too much time on their turf and they send an entire patrol after you, swarming into your tunnel as if they've been rallying their troops and waiting for the perfect moment to attack. In other biomes, we've even seen them growing crops on little farms. These are enjoyable details and occasionally make us wonder if maybe we're the monsters for invading their territory in search of loot.

### Mine crafting

The crafting is on the simpler side, and compared to most survival crafting games, there's no need to dive into wikis all that often to figure out how it all works. Farming and cooking are also pretty basic, but we mean that in a good way. Pop two ingredients in a pot – any two ingredients – and you get a meal that typically combines their properties. For instance, combining a healing berry and a flower that glows in the dark produces a meal that heals you and makes you glow. Unfortunately, inventory management quickly becomes an issue; even with max backpack slots, we found ourselves having to return to base more often than we'd like just to empty our pockets from the sheer amount of stuff being hoovered up.

As far as progression goes, there's a bit of shared DNA with *Valheim*'s boss system. A few of *Core Keeper*'s bosses are already stomping around in the world – there are even several who travel around the map – but others need to be summoned by crafting a specific item. That makes defeating them easier, because you can set up the boss chamber with traps and be as prepared as possible before the fight begins.

It doesn't always help; *Core Keeper*'s bosses are surprisingly tough, and most of them took us numerous tries and strategies before being defeated. Like *Valheim*, each main boss (there are several optional ones) gives you an item you need to progress further, encouraging



you to stray further from the safety of the core to level up your gear and weapons.

For a game that takes place entirely underground, *Core Keeper* does a great job of making you forget you're buried. In the early game, there's darkness on all sides, banished only by placing torches or hollowing out huge caverns, but there are several biomes such as the wilderness, desert and ocean that give you the feel of being outside in the warm sunlight. The lighting, effects and animations are genuinely beautiful.

It's been a while since we were so drawn in by a co-op survival game, but *Core Keeper*'s hooked us twice: once when it entered early access in 2022, and again when version 1.0 launched in August. It's got so much of what we enjoy in a survival experience: a big and beautiful world packed with danger that's exciting to explore, and a warm cosy base to come home to when we're done. **LXF**

We might complain about fishing, but then there's the farming, too...



### VERDICT

**DEVELOPER:** Pugstorm **WEB:** <https://fireshinegames.co.uk/games/core-keeper/>  
**PRICE:** £15.99

GRAPHICS	9/10	GAMEPLAY	9/10
LONGEVITY	9/10	VALUE	9/10

A sprawling survival game in a beautiful subterranean world.

» **Rating 9/10**

What is it with everyone having time to sit around and go fishing in these games?

# Roundup

Opera » Mozilla Firefox » Google Chrome  
» Microsoft Edge » Brave



**Michael Reed**

plans to retreat to a remote cabin with only a C64 and a box of games on tape for company after running five browsers for a month.

## Big-name web browsers


Browsers connect us to the world wide web, and therefore the outside world. **Michael Reed** examines five of the best for the Linux desktop.

### HOW WE TESTED...

To begin with, we made sure that all typical installation scenarios were well covered by visiting the respective project web pages and searching repositories such as Flathub and Snap. We installed the corresponding Android app for each of the browsers on a mobile phone for testing.

When not deliberately exploring the features of each browser, we spent time with each of them in an informal context, using them for typical day-to-day surfing. This meant signing into online accounts, importing bookmarks and genuinely using each one for some everyday work, such as creating this article using *Google Docs*, and tasks such as responding to emails.

As four of the five browsers use the same rendering engine, we didn't notice much difference in actual web page rendering. The outlier with its own engine, *Firefox*, is adherent to the same web standards as the Chromium engine. So, once again, it didn't perform noticeably differently in the business of actually rendering the pages.



**I**t sometimes feels as though the web browser is the most important application that we run on our Linux desktop. It's the conduit that connects us to the web and therefore the outside world. Sometimes, we're hanging out on social media, and other times we're looking around online for something specific that we need. Even further, what we do with our web browsers these days has moved beyond simply browsing the web with them, because many stalwarts of the desktop application scene have made their way over to web-based equivalents.

Of the five options we've looked at here, they are all applications that we could recommend, as these are the cream of the current crop of desktop browsers. There's always going to be an element of personal taste when choosing the web browser that is best for you, but we're considering things such as the user interface, how extensible and configurable the browsers are, and how powerful they are in terms of privacy and ad blocking. We like reliability in the basic features, but we've decided to score highly for any novel features and innovations that an individual browser offers.

CREDIT: Getty Images/ E+/Mediterranean



# Installation

We want a smooth installation experience with lots of options.

**W**e found up-to-date *Brave* packages on Snap and Flatpak, but its website informed us that these don't work as well as the native packages, which are available for most popular DEB and RPM based distros. The instructions involve cutting and pasting a long command line that adds a special repository. As the project is open source, if you have an unusual setup, you could probably install *Brave* manually.

*Google Chrome* isn't open source, despite using open source components. *Chromium* is the open source version, but it lacks some features, such as syncing with the Google cloud. For this *Roundup*, we are examining the proprietary *Google Chrome*. The website offers DEB and RPM packages, and it's also on Snap and Flathub. In our experience, occasionally, there is a muddle up with *Chrome*'s internal update system, and you have to manually fetch the package to perform an upgrade to the latest version. *Chrome* isn't available for the Raspberry Pi, but *Chromium* is.

*Edge* and *Opera* are proprietary browsers and have a few things in common. Their respective websites offer RPM and DEB packages. This might be restrictive, but Flatpak versions are also available in both cases, expanding the support for different distros. *Edge* was the only one of the browsers that couldn't be found in the Snap repository. Neither *Edge* nor *Opera* support the Raspberry Pi. When installing the *Opera* DEB package, the script offers to add the update repo to your system. However,



Thanks to being open source, Firefox is available for most combinations of platform and operating system. Beyond that, it's available through Flatpak and Snap.

some web videos didn't play properly in *Opera* until we manually installed a version of libffmpeg. Annoying!

There are a lot of reasons for using a completely open source app like *Firefox*. All the internal workings are available to see, and there is some assurance the app's not sending out data that you wouldn't like it to. It also means that there is an incredible range of installation options and platform support. You can install it via RPM or DEB, or its internal, highly reliable update mechanism.

## VERDICT

OPERA	7/10	MICROSOFT EDGE	6/10
MOZILLA FIREFOX	10/10	BRAVE	8/10
GOOGLE CHROME	7/10		

**Firefox has the widest range of platform support and installation options due to its open source nature. Chrome has an open source variant, Chromium.**

# Configuration and settings

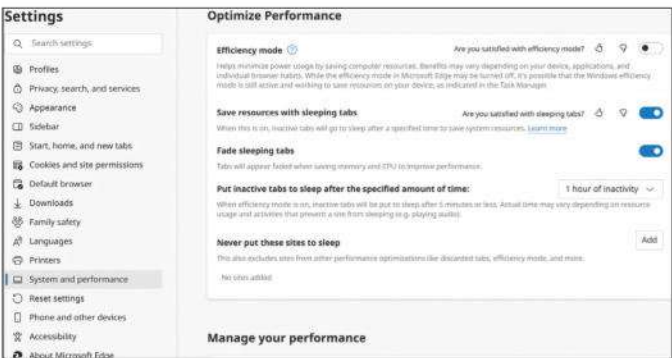
With a new browser, the first thing to do is change a few config details.

**O**pera's quick settings sidebar is a good idea, but you have to scroll to the bottom of that to get to the full settings. Additionally, we're not sure about the choices of quick settings. For example, how often will you have to toggle the settings for showing news or weather on the blank tab page or AI services in the sidebar?

When you do reach *Opera*'s full settings, there is a lack of categories and the settings are largely presented as a single list that you have to scroll down. The settings are searchable, but that can never be the full answer, as you don't always know the right term to search for. A clearout of the cluttered quick settings and more organisation in the main settings would be a good idea.

*Chrome* does a better job because there are about 15 settings categories. *Brave*'s settings page doesn't have as many categories as *Chrome*, but we couldn't decide which is the better arrangement. *Brave* had a nicer overall visual layout, and the options are much the same. Both provided brief descriptions under the options.

*Firefox* keeps the number of main categories down to around five, which feels a little understuffed. For example, the General



Like most areas of *Edge*, the Settings interface is attractive and well organised, with just the right number of categories. Where needed, descriptions are supplied.

category contains the appearance subcategory, along with subcategories for things such as language, download directory and network settings, which aren't closely related. It's not unusable, but we'd like to see more categories.

*Edge* seemed to hit the best possible balance, with quite a lot of categories and subcategories, a search bar and good overall presentation.

## VERDICT

OPERA	6/10	MICROSOFT EDGE	9/10
MOZILLA FIREFOX	7/10	BRAVE	8/10
GOOGLE CHROME	8/10		

**Edge has the best layout in its main settings, but Chrome is well organised in that department as well.**

# User interface design and feel

The main point of contact.

The user interface of a web browser is a special case because it is often a user interface on top of another one, that of the website or web application.

There's no accounting for taste. Microsoft applications like *Edge* have user interfaces that are adaptive to the task at hand. At the other end of the spectrum, the *Firefox* interface is old school because it has a more traditional interface where icons and menus are the same regardless of what context you are working within.

While there are differences in approach, the overall layout of a web browser has become fairly standardised over the years. Scoring is difficult because not every extra feature or innovation is a good idea, and some things simply get in the way. We're going to second-guess the mood of our readers: it's about time that browsers offer a multi-function sidebar, but some of the 'intelligent' features just get in the way.

Opera9/10

Mozilla Firefox7/10

The *Opera* user interface is attractive, even though it's a bit busy due to all of the extras that this browser has built in. Visual touches, such as shadows, add contrast and space between the UI elements.

A lack of mouse wheel tab flipping is a shame, but we like its tab islands. *Opera* tries to guess which tabs are related and places a coloured divider in front of that set. You can add to an island, and you can move or collapse it.

The icon-driven sidebar is extremely useful. It's convenient to be able to do a quick check of your social media feeds or the AI chat with the pop-out sidebar rather than having to change tabs and lose the flow of what you were doing.

Some of the features feel glitchy. For example, if you highlight text, a pop-up menu appears, but this can mean that a right-click is initially ignored and requires a second attempt. All of these extras can be disabled if you don't like them, though.



The *Firefox* interface is a traditional one that is starting to look a bit drab. It could sit comfortably in a government office, a classroom or a similar environment, and there are few bells and whistles. However, on the positive side, such a plain interface is easy to understand and use.

Beyond the basics of tab management it offers, *Firefox* does have an icon that takes you to a full-page overview of your open tags, recently accessed tabs and tabs on other devices, and all of it is searchable. Leave it to Mozilla to add a sensible (although not particularly glamorous looking) feature like that. It's useful.

It's a bog-standard interface, and we do wonder whether *Firefox* has started to fall behind in innovation in this area. On the other hand, you could easily argue that it's a starting point that covers the basic features and can be customised and expanded using add-ons.



# Privacy and security

Extra features to put you in control of your online safety.

*Opera* has a free, built-in VPN (virtual private network), which is nice to have on tap, even if it's arguable whether it's truly a full VPN. It does allow you to change your apparent location and it does provide extra security when enabled with a click of an icon in the title bar area. It's nice to have as an integrated freebie, even if the features aren't on a par with commercial VPNs.

*Brave* sells itself as the most privacy-focused major browser. Like most browsers these days, it makes an effort to block harmful or suspicious cookies, but *Brave* can show you what it's blocking and allows the user to fine-tune the settings on the fly for an individual site. It also has an icon to open the current website on the Tor network for enhanced privacy, which we consider to be roughly the equal of *Opera*'s VPN feature.

*Edge* has sensible configuration options for tweaking privacy and security in areas such as blocking trackers and general web browsing. In each, there are settings such as Balanced and Strict, with full bullet point summaries of what each option involves,

and it's probably the right balance of configuration options for the average user.

*Firefox* has a similar set of settings for tracking preferences labelled Standard, Strict and Custom, and some individual settings as well.

*Chrome* has a full range of privacy settings spread over multiple pages, with full textual explanations of what the functions do. The downside is that it's not set up for quickly tweaking settings and adapting to individual sites. It's rather verbose and the options are spread between multiple pages.

VERDICT

OPERA	8/10	MICROSOFT EDGE	8/10
MOZILLA FIREFOX	7/10	BRAVE	10/10
GOOGLE CHROME	6/10		

Brave and Opera promote themselves as offering enhanced security. Edge has settings that are easy to understand.



Google Chrome7/10

Microsoft Edge8/10


Brave8/10

In a word, the *Chrome* interface is slick. Google has endeavoured to remove everything that gets in the way while surfing in a maximised browser window. This means that the user interface is neither cluttered or too plain looking.

The right-click menu is fairly standard, but across different machines, we noticed a tendency to become collapsed and scrollable. We'd prefer it if it's either intelligent or traditional, but non-standard and needing to be scrolled? No.

You can create colour-coded tab groups, but it doesn't group automatically. Being able to right-click on images and start an image search on the web is handy.

There is a sidebar, but it only pops out for things like image searching and bookmarks, and we'd prefer to see it more fully utilised, as in some other browsers. *Chrome* has evolved into the standard interface for accessing the web, but we'd like to see a few extras and new ideas.



*Edge* has a pop-out sidebar for access to tools such as the AI chat, along with Microsoft-centric web-based tools such as *Microsoft 365* and *Outlook*. In fairness, there are scores of installable applets for non-Microsoft services such as Facebook and Twitch that can be added manually.

Like *Opera*, when you highlight text, a pop-up appears with common functions, which takes some getting used to. *Edge* has a side-by-side split-screen view, which ought to become standard in all browsers.

Tab management has support for groups, but they're not automatic. However, *Edge* had a high hit rate when guessing tab group categories. Switching to vertical tabs, which take up the left-hand side (a feature *Brave* also has), is an interesting option that could grow on us.

*Edge* gives us a peek at an attractive, modern user interface, but some Linux users might prefer something traditional and consistent.



The *Brave* interface looks fairly similar to that of *Chromium*, the browser upon which it is based. One upside of this is that it looks reasonably modern, but it still works in the standard way that we have come to expect. For example, if you highlight text, the right-click menu is pretty much what you'll be used to rather than a pop-up that tries to intelligently guess what you want to do next. Frankly, this is what we prefer. It's not slavishly identical to *Chromium*, and a few icons have been moved around, and things such as access to **Downloads** or adding to favourites are a bit more prominent.

*Brave* has a pop-out sidebar, which is where you access Leo, the *Brave* AI chat, and it's also a good way to access your bookmark collection in full. It's good to see a sidebar, but there doesn't seem to be much of an app selection, like the world of interesting services you can add in *Edge* and *Opera*.



# Ad blocking

We see the need for adverts, but how much intrusion is too much?

**A**dverts can present a moral dilemma because some content creators rely on us seeing them in order to be able to provide our favourite content. On the flipside of that argument, many adverts are intrusive to the point of ruining the entire web-surfing experience. We like a nuanced solution, where you can make decisions about the type of advert that gets through.

*Opera*'s ad blocker is built in and enabled from the beginning. By default, it's set to block "intrusive" ads. In use, occasionally, we'd see an advert for a company such as Disney, but we didn't get pummelled by anything too garish.

Like its privacy features, *Brave* has comprehensive ad-blocking options that are enabled by default and highly configurable. Selecting between "acceptable" adverts (the default) or no adverts, or disabling the ad blocker for the current website can be done quickly from within the interface. The settings can allow ads on some social media sites, and there's a novel system where you can earn Basic Attention Tokens for

viewing certain ads, and these tokens can be exchanged for actual currency.

*Chrome*, *Firefox* and *Edge* don't come with ad blockers built in, but they have ad-blocking add-ons, so you can decide how you want to approach the subject. The most popular plugins tend to default to a blanket approach to blocking all ads.

*Edge* has adverts on its own startup page that the ad blocker didn't seem to be able to do anything about. *Brave* sells advertising space within its browser. While testing, we only noticed these ads in the new tab news feed.

VERDICT

OPERA	8/10	MICROSOFT EDGE	5/10
MOZILLA FIREFOX	5/10	BRAVE	7/10
GOOGLE CHROME	5/10		

Brave and Opera come with ad blocking enabled out of the box and good configuration options.

# Mobile app

We like to carry on browsing on the move.

All the browsers supported the major desktop operating systems along with Android and iOS mobile OSes, and they supported syncing bookmarks and other data between them. We installed the Android apps to get an idea of what the experience was like and to make sure there was a good workflow when moving from one system to another.

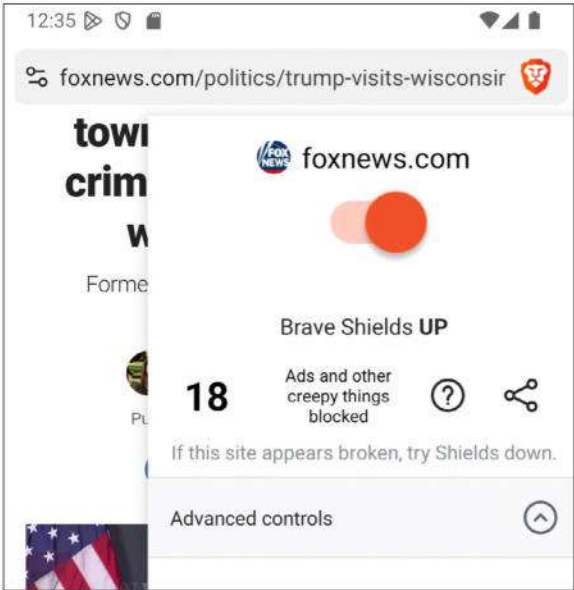
The most important feature, synchronisation of bookmarks and login data, worked for all five apps, and is also supported by the desktop browsers between platforms.

When it came to pairing the phone app with the web browser accounts, *Firefox* and *Opera* worked smoothly, as we simply scanned a QR code on the screen from the appropriate website. We used email/password logins on the other browsers.

All the mobile browser apps have settled into a similar style, with a homepage featuring a search bar and icon links to favourite sites. Below this, there is typically a scrollable news feed. It's perfect for light reading, and the *Opera* app can make that reading even lighter with a built-in AI summariser function. We also liked *Opera*'s Flow system for sharing notes and links with your other devices, but all of the apps can do some sharing.

*Brave*'s Shield interface is prominent in the mobile app, giving the same fine-grained control as the desktop browser as to what trackers and ads are being blocked.

As usual, the *Firefox* app delivers the basics very well. The app isn't quite as extensible as the desktop browser, so if there is something you need in the app, you might not be able to add it.



Brave's Android mobile app. This carried across logins, bookmarks and favourites as expected. It also carried across Brave's interface for controlling the ad and tracker blocking facilities.

VERDICT			
OPERA	8/10	MICROSOFT EDGE	8/10
MOZILLA FIREFOX	6/10	BRAVE	9/10
GOOGLE CHROME	7/10		

A fairly standard mobile browser interface seems to have emerged that they all adhere to. Brave, Edge and Opera add extra features, though.

# Extra features and add-ons

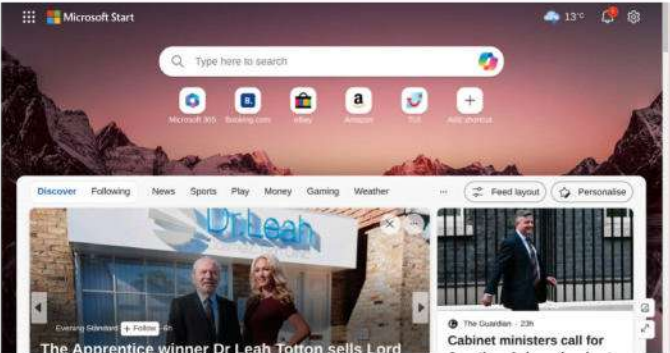
Extra functions that make it stand out, and potential for extendability.

*Opera*, *Brave* and *Edge* integrate AI assistance. In each case, the AI window can pop out of a sidebar for quick access. This makes sense as you're typically trying to get to the bottom of a subject that is dominating the current window. As with all AI chat, quality and factual accuracy is a roll of the dice.

When playing video, *Opera* can activate its Lucid mode and enhance the video quality. Flow is another good example of the type of extra feature that *Opera* throws in. It's for sharing notes, links and files between devices while you're browsing.

*Chrome* and *Firefox* both have hundreds of extensions created specifically for each – though *Chrome* has started to lock down on these. So, if there is a feature that you need from another browser, you can almost certainly add it. For example, neither makes much use of its sidebar but there are many sidebar add-ons. On the downside, you have to sort through all the extensions, add and configure them.

*Brave*, *Opera* and *Edge* have their own extensions, but as they are partially based on *Chromium*, many *Chrome* extensions work within them. The rub is that there is no guarantee that a *Chrome* extension will work fully or predictably in those browsers.



Edge's new tab page is well laid out. It's a shame that browsers such as Chrome and Firefox haven't started to capitalise on adding features to the blank tab page yet.

*Brave*'s new tab page is a sort of dashboard with inspiring photography as a backdrop, icons linking to frequently accessed pages and some statistics. *Edge* makes a similar effort, but *Opera* strips things down to site icons and search, and much the same can be said for *Firefox*. *Chrome* features only search. It might seem like a small feature, but *Edge* pipes in some spectacular landscape photography where appropriate.

VERDICT			
OPERA	9/10	MICROSOFT EDGE	9/10
MOZILLA FIREFOX	6/10	BRAVE	9/10
GOOGLE CHROME	5/10		

Chrome and Firefox provide basic features but can be expanded through add-ons. Edge, Brave and Opera have enhanced features out of the box.



# The verdict

## Big-name web browsers

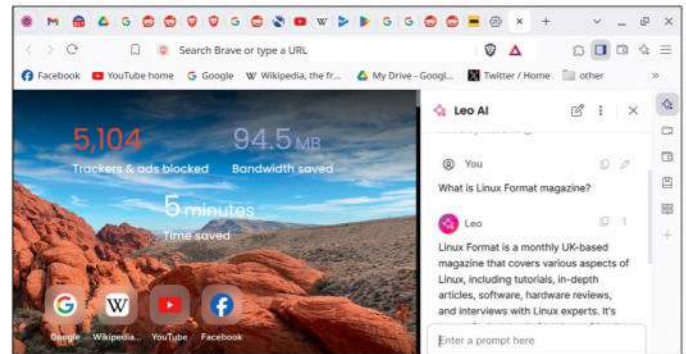
**B**rave feels like a remix of *Chrome* but with some welcome additions, such as AI integration, an extensible sidebar, and detailed security and privacy features. The additions to the interface, such as the sidebar and right-click-accessible AI functions, are genuinely useful. These features are great, but it doesn't feel like it ever crosses the line with interface elements that get in the way. It's also completely open source.

*Opera* uses *Chromium* as its base these days, but it's certainly still a browser in its own right. The developers have probably made the right decision on this matter rather than creating every bit of code independently, like they used to. How much you like *Opera* might depend on where you stand on how bare or pre-customised you like your browsers. If there is an *Opera* UI refinement that you simply can't get on with, you can probably disable it. It's packed with premium features and extras that you'll discover as you get to know the browser through use.

*Chrome* is the number-one browser in the world right now. It's standards compliant, pleasant to use, and reasonably fast. It's been around for nearly a decade and a half, and the user interface could use a few additions, such as an icon-driven sidebar and automatic grouping of tabs. You'd think that a company as big as Google could add a few innovations to improve the overall web-browsing experience. Google has recently announced a change to how extensions work that could invalidate some of the most popular ones. By comparison, a vote for *Firefox* is a vote for long-term stability.

We like *Firefox* very much, but it's fallen behind when it comes to the newest features. Many will like a plain browser, and its saving grace, in this regard, is that you can add anything you can think of by using its extensive add-ons system. But this means you have to do that customisation work yourself.

Even though it uses some *Chromium* code for the web layout engine, *Edge* is every inch a Microsoft product. The user interface is likely to divide Linux users because some elements pop up or hide themselves according to the current context. Even though it's not open source, the add-ons ecosystem is extensive. Overall, it's a perfectly fine browser even though it has the expected integrations with Microsoft services.



### 1st **Brave**

**9/10**

**Web:** <https://brave.com>

**Licence:** Mozilla Public Licence 2.0 **Version:** 1.70.119

Privacy focused. Covers basics with some enhanced features. Open source.

### 2nd **Opera**

**8/10**

**Web:** [www.opera.com](http://www.opera.com)

**Licence:** Proprietary freeware **Version:** 114.0.5282.21

A premium feel with extra, useful, modern features. Beautiful user interface.

### 3rd **Mozilla Firefox**

**8/10**

**Web:** [www.mozilla.org/en-GB/firefox](http://www.mozilla.org/en-GB/firefox)

**Licence:** Mozilla Public Licence 2.0 **Version:** 130.0.1

Basic browser that does all it should. Highly configurable and extendable.

### 4th **Google Chrome**

**7/10**

**Web:** [www.google.com/intl/en\\_uk/chrome](http://www.google.com/intl/en_uk/chrome)

**Licence:** Proprietary freeware **Version:** 129.0.6668.70

All the basic features with a slick interface. Highly extendable with add-ons.

### 5th **Edge**

**7/10**

**Web:** [www.microsoft.com/en-us/edge](http://www.microsoft.com/en-us/edge)

**Licence:** Proprietary freeware **Version:** 129.0.2792.65

Good-looking interface layout with lots of extra functions. Feature packed.

## » ALSO CONSIDER

*Seamonkey* is based on the same browser technology as the *Firefox* web browser, but it's an entire internet suite that includes an email client, an HTML editor and an IRC client in addition to the browser itself. Internet suites are great in cases where you're setting something up for other people to use and you want to standardise the whole experience.

*Vivaldi* tends to be thought of as a feature-packed browser for experts who don't mind navigating a lot of advanced

features. We skipped over it this time, however, because it covered a lot of the same ground as web browsers such as *Brave*, *Edge* and *Opera*, and we wanted to assess *Firefox* and *Chrome*, as the two most popular Linux browsers, at the same time.

There are plenty of smaller or niche browsers, such as text-mode browsers and solid choices like *Konqueror*, but we wanted to concentrate on the heavyweights this time. **LXF**



# POWER UP FEDORA!

Not all heroes wear capes. **Jonni Bidwell** has got a brand new Fedora to discuss and these hat jokes are un-brim-leavable.

**T**he mid-autumn distro season is always fun. We get a new interim release of Ubuntu (see page 20) and a brand new version of Fedora. Both use the Gnome 47 desktop, so on the surface there's not a lot to separate them. We're often accused of favouring Ubuntu, so let's counter that. Fedora aficionados won't need any reason to switch. If you're coming from another distro, though, you might be curious about Fedora. Perhaps even trepidatious. Well, we're here to tell you that Fedora is for you. Never mind what you might have heard, it's really no more advanced than Ubuntu, or any distro based on it. We'll show you how to install it, get to grips with Gnome, and tailor Fedora to your whim.

A raft of software is available on Fedora. And that raft has got much bigger thanks to the option of easily adding third-party repos. So, it's easy to get proprietary software, and you can get the latest tools straight from the developers thanks to Flatpaks. Now you can potentially get software that's newer than those annoying Arch users have. You'll be a Gnome keyboard shortcut whizz thanks to our guide. And if that's not your cup of tea, we'll show you how to install all manner of alternative desktops.





# Fedora is for all

You don't need to be some sort of geek to don a Fedora.

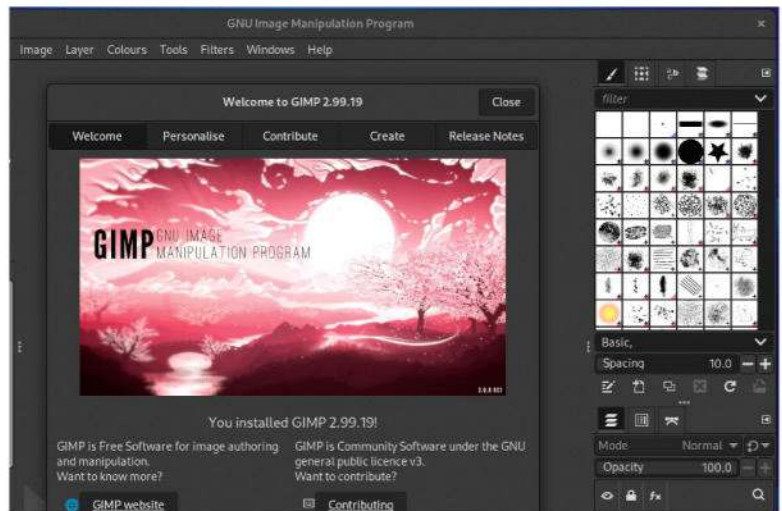
**F**edora Linux has always been something of a trailblazer. In one sense it guides CentOS Stream, which in turn sets the road map for the commercial Red Hat Enterprise Linux (RHEL). From a desktop point of view, it's often the first (ahem, Arch BTW) to adopt new features (such as Wayland, Btrfs, Vulkan). Thanks to its (roughly) six-month release cadence, it equals the interim Ubuntu releases, and is way ahead of Ubuntu's two-yearly Long Term Support (LTS) offerings as concerns new software. Updating your OS twice yearly sounds like a pain, and indeed it used to be. Partly for that reason, partly due to the extra hoops required to install proprietary software, and partly because earlier Fedora releases didn't quite have Ubuntu's level of user-friendliness, Fedora has often been put in the 'intermediate distribution' rubric.

That's just not true any more – the OS updates are seamless, and Fedora is every bit as user-friendly as modern Ubuntu. Indeed, in an era when the Ubuntu team is taking a lot of flack for its design decisions (in reverse chronological order: Snaps, diagnostic data, replacing the Unity desktop with a Franken-Gnome, and partnering with Amazon), Fedora is enjoying widespread recommendation as a beginner's or gamer's distribution. The matter of proprietary software (which once involved manually adding a third-party repository) is now solved by clicking a simple tickbox in the installer. This one click gives you access to the repos for the Nvidia graphics driver, *Google Chrome*, and the *PyCharm* integrated development environment (IDE) for Python.

## No place like Gnome

Fedora's flagship release offers the most faithful experience of the Gnome desktop you can find. So, (without the help of extensions) you won't get a Dock at the bottom of your screen, a system tray, or Ubuntu-style application indicators. This 'Gnome as the Gnome developers intended' look makes Gnome seem minimal or perhaps hyper-modern. And that seems to be something users either love or hate. We'll show you the handful of keyboard shortcuts and everything else you need to embrace the Gnome workflow over the page.

You'll still see some criticism – Gnome in general is a bloated desktop, but in an era when Windows installs want 64GB of space and almost as much memory, that criticism seems less valid. Gnome is a fully featured desktop environment, complete with



integrated document searching (which requires a big database), accelerated rendering (requiring a graphics card less than about 15 years old), and an easy dialog for disabling optional services. It will almost certainly run on any hardware that could run Windows 7. But for old machines, it will be a challenge, and there are more suitable desktops (such as LXQt and Xfce) and distros (such as Bodhi or Lubuntu) out there.

Fedora makes it easy to install the Nvidia driver, and has done so for some time. This time, though, we're spared another headache because the package now supports Secure Boot. So, you won't have to (as in previous releases) either disable Secure Boot or manually sign the module and enrol your own MOK keys. Since the proprietary Nvidia driver isn't included on the live medium, and in the unlikely event your card requires it, there's an option to start with basic graphics in the Troubleshooting boot menu.

GIMP "3.0" is available in Fedora's repo, even though it hasn't been released as we write this.

## » FEDORA GIVETH AND FEDORA TAKETH AWAY

For a long time, Fedora has been pushing Wayland (as opposed to the ageing X.org display server) as the default graphics protocol in Gnome. With Fedora 41, the inexorable march towards the retirement of X (the graphics thing, not the hate pool that used to be a reasonable social media platform) takes another step as X.org is banished from the default install. It is still available in the repositories, and we'll show you how to install it later. It's worth noting that there is still an X server installed, in the form of Xwayland, the translation back-end that allows old apps that don't support Wayland to run.

The now commercially licensed *Redis* database has also been replaced by the FOSS-friendly *Valkey*. Python 2 is finally gone in this release. Like the parrot in the *Monty Python* sketch. One of the few reasons it was kept around was that it was required by *GIMP*. But in September, it was announced that *GIMP* 3, which is finally free of the obsolete language, would be included in Fedora 41. We've been banging on about migrating to Python 3 since at least **LXF195**, so hopefully any packages you use have also upgraded. Otherwise they, too, will have found themselves deprecated.



# Popping on Fedora

Grab yourself a cup of tea and a USB. It's time to fortify and update your system with Fedora (and tea).



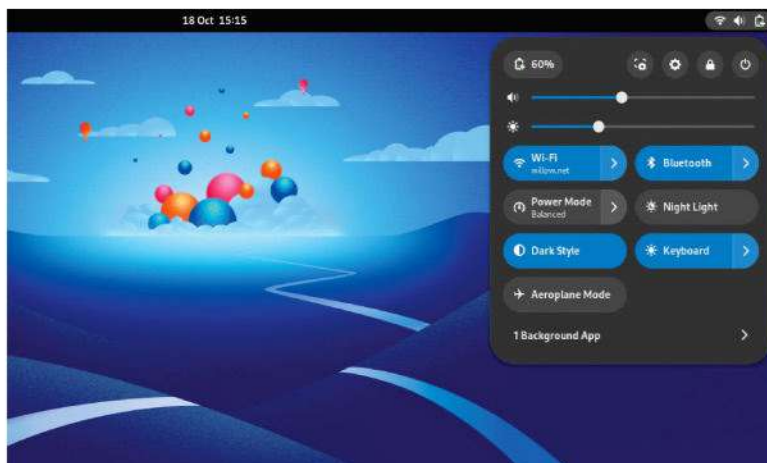
## QUICK TIP

Download the Linux Etcher Appimage from <https://etcher.io/>. Right-click on this and select Properties > Permissions > Allow executing of file, you can now run it. To use the Terminal run `cd ~/Downloads and use chmod +x etcher<version>.appimage and then run it.`

**I**nstalling Fedora is a breeze. First, grab the ISO from <https://getfedora.org>.

It's only a couple of gigabytes, so hopefully won't take long to fetch. Best make a cup of tea anyway, just in case. Next you'll need a USB stick (4GB or higher) and a tool for writing to it. We always recommend *Etcher* (<https://etcher.io>) because it's cross-platform and pretty much self-explanatory to use, see the top tip left for more details. Once it's written, you need to tell your BIOS or UEFI to boot from the USB stick. You might be able to get a pop-up boot menu with F12 or F10, or you may have to hit Del or F2 (or any other key combo) to enter the interface and configure the boot order from there. If you get stuck booting into your previous OS, you can access UEFI settings from Windows by holding down the Shift key while pressing the Shutdown button in the Start menu. In Linux, you can run `sudo systemctl reboot --firmware-interface` to force the issue.

Once the desktop is loaded, you'll see some shortcut icons in the Dock at the bottom, an overview of the desktop (with the welcome screen), and a fraction of another desktop to the right. The latter isn't some graphical aberration, but it's where you'd click if you wanted to use a new virtual desktop. More on that



An unblemished Gnome desktop, with the status menu open for easy access to common dials.

later. This view as a whole is Gnome's Overview mode, and can be summoned at any time with the Windows, or Super to use the OS-agnostic term, key. Click on the desktop and we'll get on with installation.

If you want to explore Fedora risk-free, though, click Not Now and you can explore the live environment. This is very similar to how Fedora will look and act once installed, albeit a tad slower due to it running off USB. Go back to the Overview with the Super key, and click the grid of dots at the right of the Dash. This shows all the apps currently installed. There are a few more in a real installation, but you'll see the core Gnome applications as well as the *LibreOffice* suite. You can search for an app by typing the first few characters (no need to click anywhere) and then hitting Enter once it's highlighted. When you're ready, fire up the installer.

## » KEYBOARD SHORTCUTS

We mentioned that your Gnome workflow will be made much smoother if you grok a couple of keyboard shortcuts. So, without further ado, let's avail you of a few. We've already covered the Super key to toggle the Overview, and you'll already be familiar with Alt+Tab for switching between applications. Here are some others:

<b>Super+A</b> .....	Show all applications
<b>Super+M</b> .....	Show calendar and notifications
<b>Super+Left/Right</b> .....	Tile window on left or right (split view)
<b>Super+Up/Down</b> .....	Maximise/restore a window
<b>Super+Page Up/Down</b> .....	Switch virtual desktops
<b>Super+Shift+Page Up/Down</b> .....	Move window to another desktop

If you DuckDuckGo for "gnome keyboard shortcuts", you'll find a handy cheat sheet with the lot.

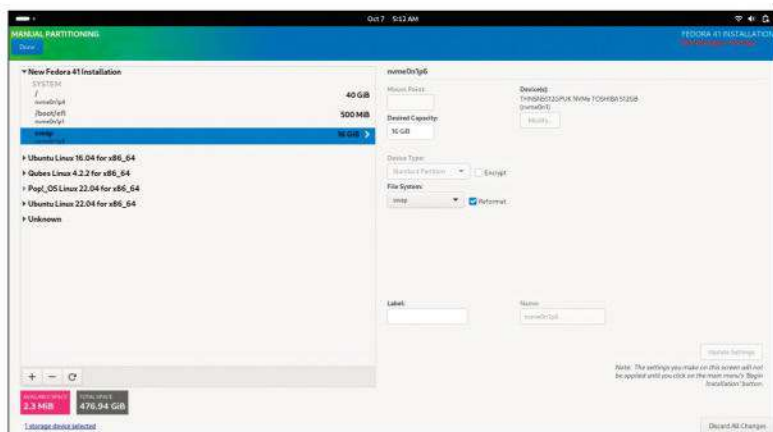
Oh, we almost forgot – Gnome has a built-in screen recorder, too. It involves a bit of finger gymnastics (Ctrl+Alt+Shift+R) but it's a great tool for making postmodern desktop videos.

## Perfect partitioning

Once you've chosen your locale settings, you're presented with an installation summary. You'll see an exclamation mark in the Installation Destination section, so click that. Automatic Installation is easiest, and if you already have an empty disk, it's only a matter of a single click to set up. If there's data on the disk where you want Fedora to live, and you're sure you don't want it, click the Reclaim Space box at the bottom-left. This lets you delete whatever partitions you don't want. If you have more advanced partitioning needs, use either of the custom options. Our initial install occupied just over 7GB, but ideally you'd want upwards of 40GB for your root partition.

Fedora's default Btrfs partition scheme will be better than fine for most people. Btrfs is a next-generation filesystem that is designed to be fast and robust. It has built-in snapshot and RAID capabilities, as well as checksumming to mitigate against silent data





It's easy to install alongside other distros; the minimum required is specifying root and EFI partitions.

corruption. You only get that for free, though, if Fedora is allowed its own drive. Otherwise partitions can be configured manually using either *gparted* or Fedora's own *blivet* tool.

When you're happy with your arrangement, click Done at the top-left (nothing gets written or deleted at this stage). Then click Begin Installation, whereupon the installation begins and there's no going back. Once it's done, click Finish Installation to return to the live environment. When you're ready to see if the install worked, click the status icons in the top-right, then click the power icon (top-right of the pop-up menu), and then click Restart. Fingers crossed, you're greeted by another welcome screen and invited to start the setup process. Don't worry – it only takes two ticks. You're asked to set a couple of privacy options, enable the third-party repositories we mentioned earlier, and set up a username and password. If you are setting up a business machine, you can also configure an Active Directory domain (just hit the Enterprise login button).

The Welcome tour introduces you to the Gnome desktop, starting with the Overview, then working through some three-finger touchpad/touchscreen gestures: triple-finger swipe left/right to switch virtual desktops, and tri-digit swipe up to bring up the Overview. Simple. It's worth reiterating that you can search not just applications, but documents from your **Home** folder, contacts, software and more. You can even run in-line calculations; try typing `2 + 2` into the Overview. Other Gnome apps can add their own search facility to the Overview, making it even more all-powerful.

One of the most talked about features in Gnome 47 is the option to add accent colours. This might not sound like much, but it is a surprisingly big change. For a long time, Gnome 3 (without any extensions) hasn't had any provision for custom theming, beyond backgrounds and light/dark themes. Now you can choose from nine (count 'em) colours. If you don't like any of them, you can set your own with third-party tooling, which we'll get to. In the meantime,

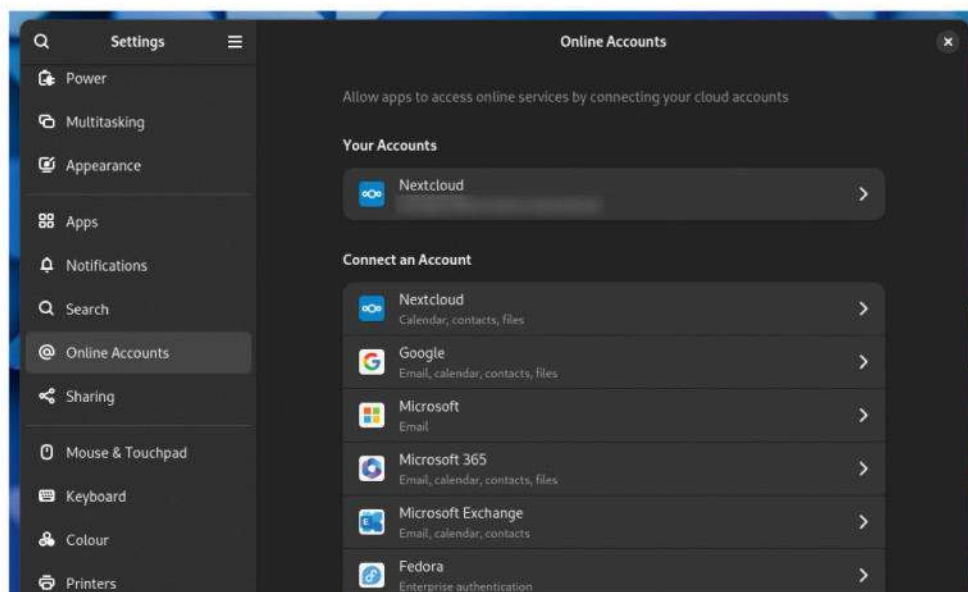
access these new accent options from the *Settings* app. Open the status menu in the top-right, and click the cog. Or open the Overview (with the Super key) and type the first few letters of **settings**. You'll find the colourful options in the Appearance tab.

## Super settings

The Settings panel is a one-stop shop for system configuration. If you don't like the idea of Gnome searching and indexing the files in your **Home** directory (and sharing them among other Gnome apps), you can turn this off, or restrict

what gets indexed, from the Search tab. If you prefer natural touchpad or scrollwheel scrolling, this and other options are available from the Mouse And Touchpad area. If you're privacy conscious, you can turn off automatic problem reporting (anonymised) and configure location services from the Privacy And Security tab. Gnome can integrate with all manner of online services (Nextcloud, Google, Microsoft 365 and Exchange, and more), so you can access your cloud storage, email, contacts or whatever from the relevant Gnome app. The rest of the Settings page is largely self explanatory, so we'll leave you to figure it out.

Multi-monitor (and HiDPI) support continues to get better and better in Gnome. If you're the kind of person who unplugs an extra monitor without first moving applications off it (you savage, you), Gnome has your back. In this situation, applications are not banished off the edge of a remaining screen, but are gracefully located to another display. If you were to plug back in that same monitor, applications are automatically returned there, as though nothing had happened. Fractional scaling has gotten even better in Gnome 47, and should be available without further action for HiDPI displays. In versions past, you'd have to manually enable it using a GSettings key, or face blurry text when running apps through Xwayland.



Online accounts are simple to set up in Gnome. All your calendars, contacts and cloudy paraphernalia can be integrated.




# Getting around Fedora

Now that we've got to grips with the desktop, let's look at how we install software. And mess with the desktop.

**S**ooner or later, you're going to want to install some extra software. You can do this from the command line, but it's much easier to use the app. Which is called *Software*. There's a shortcut to it on the Dash (it's the white briefcase emblazoned with some shapes). Open it up and you'll see some featured applications. Scroll down to see more, or explore some of the categories at the top. We like to keep some files on our venerable Nextcloud server up in the clouds, so we can access them from anywhere, so we'll walk you through installing the desktop client; feel free to install something else. Search for Nextcloud from the top-left of the *Software* app and click the Nextcloud Desktop Client entry. If you enabled third-party repositories (in particular the Flathub repo), you have a choice of where to install it from. Fedora's traditional RPM repos are pretty up to date, but in general, if an app is available on Flathub, too, a newer version is likely available there. Let's install the Nextcloud Flatpak by clicking the drop-down below the Install button and selecting the Flathub source. Now click Install and fetch a cup of tea while it installs. We'll talk more about Flatpaks when you return.

Apps on Flathub ship as modern Flatpak packages. These are self-contained affairs that make it easier for developers to ship new software without having to worry about dependencies and wait for distribution packagers to catch up. They use modern kernel features (namely namespaces and containers) to offer a degree of compartmentalisation from the rest of the system. In theory, this protects against rogue apps messing with other Flatpak apps or the rest of the system. In practice, many applications aren't much use without access to the system (for example, they want to access files in your **home** directory, rather than their own sandboxed storage). Flatpaks for fully featured apps generally depend on other Flatpaks (such as the Gnome and KDE platform runtimes). This dependency



The *Software* application is your gateway to a world of productivity or distraction. It also knows if you're using mobile broadband.

web can get quite large, so in some cases it seems like a Flatpak app is downloading a huge amount of data during installation. But this usually diminishes with time, once you've ingested some common service Flatpaks.

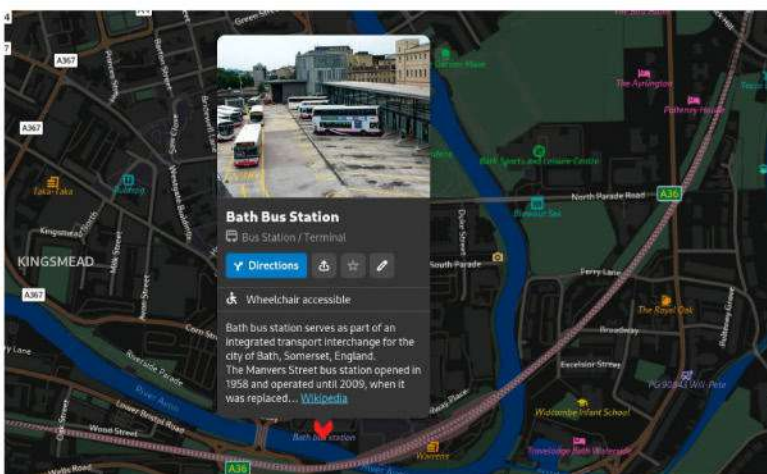
For all their newness and goodness, you might run into some theme problems with Flatpak apps. The Nextcloud client, for example, uses Qt, so doesn't look anything like the rest of the Gnome desktop. It doesn't even look much like how it would look if it wasn't a Flatpak, since no Qt themes are installed. If you have a Google, or better DuckDuckGo, around, you'll see how the *Kvantum* theme manager can help with this. Modern *GTK4* apps don't suffer from this as much (thanks to theming all being done in the Adwaita style – the only one), and *GTK3* theme Flatpaks are available on Flathub. Again, DuckDuckGo is your friend. If you want to control Flatpak permissions (without using the command line), check out the *Flatseal* tool.

The *Software* app also handles updates, as you may have gathered if you spotted the Updates button in the header bar. You're alerted to these as they become available and you'll be pleased to hear the process is much less painful (and forced upon you) than Windows updates. In some cases, a restart is required and the system boots up into an update environment that ensures everything is carried out smoothly. You can toggle between the graphical progress screen here and the textual output from the updates by pressing Esc. The hamburger menu in the top-left of the *Software* app lets you configure software sources (for example, if you have a change of heart about third-party repos). Or you can configure updates and associated notifications by clicking Preferences.

## Tweaker's corner

Gnome 3 is often criticised for lack of customisability. Most of this is unwarranted; there's a simplicity and

Gnome Maps now provides dark theme support, as well as some public transport information.





elegance that would be easily ruined by giving every widget and aspect a boatload of options. Still, for a long time, users have turned to the *Gnome Tweaks* utility to provide access to some of those hidden-away options (they could and can also be set using GSettings keys from the command line). You'll find *Gnome Tweaks* via the *Software* app, and if you install it, you're told that its extensions support has been deprecated and that you should instead install the *Gnome Extensions* Flatpak. *Tweaks* can still do a lot, though, including setting fonts and startup applications. If you really want, it can even put the maximise and minimise buttons back on window title bars.

Extensions can change the whole way Gnome looks and operates. But they're disabled by default for performance and stability reasons. You might want to throw caution to the wind and dabble with a few. In which case, search *Software* for the third-party *Extension Manager* tool. We'll look at that over the page. There's also the official *Gnome Extensions* tool, but that one doesn't let you install extensions – it only lets you configure incumbent ones.

We haven't covered much that's command-line-related yet, so now's a good time to point out that the *Terminal* application in Gnome has received a huge update. Modern Linux desktops no longer require the

```

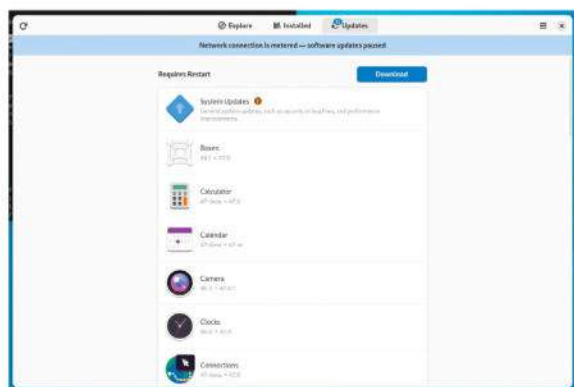
lxf@fedora:~$ sudo dnf update --refresh
[sudo] password for lxf:
Updating and loading repositories:
Fedora 41 openh264 (From Cisco) - x86_ 100% | 1.9 KiB/s | 989.0 B | 00m00s
google-chrome 100% | 5.5 KiB/s | 1.3 KiB | 00m00s
Fedora 41 - x86_64 - Test Updates 100% | 75.2 KiB/s | 17.1 KiB | 00m00s
Fedora 41 - x86_64 - Updates 100% | 49.6 KiB/s | 26.4 KiB | 00m01s
RPM Fusion for Fedora 41 - Nonfree - S 100% | 14.0 KiB/s | 7.3 KiB | 00m01s
Fedora 41 - x86_64 100% | 10.7 KiB/s | 3.7 KiB | 00m00s
RPM Fusion for Fedora 41 - Nonfree - N 100% | 37.1 KiB/s | 7.6 KiB | 00m00s
Copr repo for PyCharm owned by phracek 100% | 2.9 KiB/s | 1.8 KiB | 00m01s
google-chrome 100% | 2.4 KiB/s | 3.1 KiB | 00m01s
Fedora 41 - x86_64 - Test Updates 100% | 698.8 KiB/s | 2.8 MiB | 00m04s
Fedora 41 - x86_64 100% | 1.2 MiB/s | 3.0 MiB | 00m02s
Repositories loaded.
Package Arch Version Reposit Size
Removing:
kernel x86_64 6.11.0-0.rc5.43.fc41 anacond 0.0 B
kernel-core x86_64 6.11.0-0.rc5.43.fc41 anacond 69.5 MiB
kernel-modules x86_64 6.11.0-0.rc5.43.fc41 anacond 62.5 MiB
kernel-modules-core x86_64 6.11.0-0.rc5.43.fc41 anacond 36.7 MiB
kernel-modules-extra x86_64 6.11.0-0.rc5.43.fc41 anacond 2.7 MiB
Upgrading:
adobe-mappings-cmap noarch 20231115-1.fc41 updates 15.2 MiB

```

user to do anything involving the terminal (at least until something goes wrong). But it can be a much more efficient way of working and generally looks cool. Most prominently, *Terminal* now has a tab overview mode. Much like Gnome's Overview, this gives you previews of each session (and a similar zoom effect). A traditional tab bar appears if you have more than one tab open, so don't fret. You can also switch between (blinding) light/dark modes from the hamburger menu.

One reason you might want to use the terminal is to sample the new *DNF 5.0* package manager. *DNF* replaced the venerable *Yum* (and remains command-compatible with it) back in 2015. As an aside, *YUM* stands for *Yellow Dog Updater - Modified*, having its origins in Yellow Dog Linux, which existed from 1999-2012. In a nod to this, *DNF* stands for *Dandified Yum*. Some people shouldn't be allowed to name things. Anyhoo, *DNF* is used behind the scenes by *Software*, but if you want to see it close up, fire up *Terminal* and run `dnf -h` to see the general help page. *DNF* works, functionally, in much the same way as *Apt* on Debian (or *Ubuntu*), but it's modern and faster. *DNF* automatically updates its repository caches if they get too old. So, there's rarely any need to explicitly run `dnf update`. If you were to run, say, `dnf install quake2` (*Quake 2* was open sourced in 2023), the update would be carried out first and you'd be blasting berserkers in no time.

Gnome's new *Terminal* application turns a fetching shade of red when running stuff as root.



Don't worry, system updates are a breeze with Fedora.

## » HOW TO GET YOUR X.ORG BACK

In the unlikely event you encounter problems with Wayland (mostly, we'd wager, due to the proprietary Nvidia driver), you might consider trying your luck reinstating the X.org display server instead. Most issues with screen recording and desktop sharing have been ironed out in Wayland, but hardware configurations are a box of frogs, so you might find a bug.

Getting the X.org session back is just a matter of installing a couple of

packages. Open a terminal and run:  
`$ sudo dnf install gnome-session-xsession gnome-classic-session-xsession`

You're asked for your password because `sudo` runs the command with administrative privileges (since it's installing system files). Hit Y to confirm the operation (or 'transaction').

Now log out, and at the password screen, click the easy-to-miss cog in the bottom-right. You'll see a choice of

sessions pop up. Pick Gnome On X.org and log in to see if it works.

As we're full of surprises, we installed the X session for Gnome Classic, too. Gnome Classic adds some traditional elements (such as the Applications and Places menus, as well as a taskbar and desktop pager), while retaining Gnome's modernism. It's installed by default; we just added the Gnome Classic X session in case you want to run it on the classic display server.

# Trick out your Fedora

We look at extending Gnome with, er, extensions. Plus, a look at some of the other desktops and Fedora flavours available.

**W**e've shown you the Gnome basics, so now let's look at customising it properly (and – see *box* – circumventing it altogether). We said we were going to show you some extensions and we are people of our word. So, install *Extension Manager* from the *Software* application. Ubuntu made a reasonable concession to its users when it switched (back) to the Gnome desktop in 2018. It gave them a mouse-accessible dock, which is provided by a Gnome extension. Many users, no matter how they try, can't live without a dock. So, let's look at installing one.

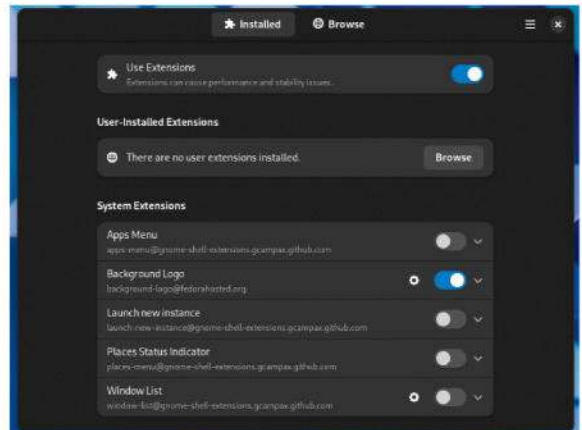
You'll find Ubuntu Dock at <http://extensions.gnome.org> along with an entry saying not to use it. Because that's not the real dock, it's just a mock extension to reserve the name. Instead, we'll install Dash to Dock, which is what Ubuntu Dock was originally based on. As its name suggests, it takes the Dashboard out of Gnome's Overview and puts it on the desktop, making it a Dock. This nomenclature is getting complicated. Open *Extension Manager* from the Overview and once it loads you'll see (in the System Extensions) that Fedora, too, has some Gnome extensions installed behind the scenes (only the one for overlaying the logo on the default desktop is enabled, so the experience is still very vanilla Gnome).

## OUT OF THE BLUE

“Silverblue (and its atomic siblings) separate the system image (the immutable part) from any apps and user data”

Go to the Browse tab at the top. You're presented with a randomised list of extensions, so have a gander at what's on offer. Then search for “dash to dock” at the top and click Install. It won't be enabled by default, so go back to the Installed tab and activate it. Hey presto – a Dash should appear at the bottom. It will be hidden if there's no room (for example, if a window is maximised), so move the pointer to the bottom of the screen and it should pop up. Bringing a traditional metaphor back to this desktop dystopia.

You can configure the Dock by clicking the cog next to it in *Extension Manager*. It doesn't interfere with the Dash in the Overview, because they are in fact one and the same thing. So, if you move the Dock to the left of the screen (like Ubuntu used to have way back when), it stays exactly there when viewed from the Overview. The freshly inaugurated Dock can be



The Extension Manager tool makes easy work of wrangling your Gnome extensions.

configured in more ways than you can think of. One very useful option, in its Appearance settings, is the ability to not show the Overview immediately on login. After all, with a Dock visible from the desktop, a mouse-orientated user would launch their apps from there, cutting out the Overview middleman.

Other popular extensions include *Blur My Shell*, which adds a stylish blur to various elements of the desktop, much like Ubuntu's HUD used to do in the Unity days. If you want a traditional Applications (or Start menu), you can enable the already installed *Apps Menu* from the Gnome Classic session. If you want the Dock to instead be a fully featured Panel (with a system tray and power/status icons), check out *Dash to Panel*. That way, you can do away with the top bar altogether, and enjoy a much more orthodox desktop.

## Fedora at large

We've done a terse overview of the alternative desktop spins (and Labs) of Fedora in the box. But there's yet more to the Fedora pantheon. Like Ubuntu, there's a Server edition, which does away with the desktop and gives you *Cockpit* (Fedora's web-based remote administration tool), *Podman* (its engine for managing containers) and *WildFly* (its application server for deploying turnkey apps). Fedora has gone big on containers, and to that end there's *Fedora CoreOS* (see *LXF211*) – a minimal OS designed for running containers at scale. There's also the IoT edition, which can run on your Raspberry Pi (Workstation builds are also available for the Pi) or other single-board computers. It's container-based, too. And there's also a cloud image (also container focused) for secure deployment in the sky (or someone else's data centre).

Fedora has more desktop offerings, too. The *Fedora Atomic* project was launched in 2018 with







Installing things, such as Nextcloud, pictured here, from Flathub ensures you get the latest, and hopefully greatest, apps.

Fedora Silverblue (see LXF299) – an atomic desktop based on an immutable image. This doesn't mean it will explode on the remains of Ozymandias's statue, but rather that updates are carried out in a single transaction in a robust (and easily rollback-able) manner. Silverblue (and its atomic siblings) cleanly separate the system image (the immutable part) from any applications and user data, using RPM-ostree technology. Besides the original, Gnome-based Silverblue desktop, there are also offerings using KDE Plasma (Kinoite), Sway and Budgie.

Like any good distro, it's easy to install additional desktops (and their respective application suites) on Fedora. Newly installed desktops appear in the Session menu when you enter your password. To install KDE Plasma, for instance, open up a terminal and install the corresponding package group:

```
$ sudo dnf install @kde-desktop
```

Package groups in DNF are denoted by the @ symbol. You can see some of the other offerings by running `dnf group info`. You'll see there are groups for Enlightenment (as used by long time LXF fave Bodhi Linux), KDE Plasma Mobile and the exciting new COSMIC desktop from System76. There are also groups for the various spins available (see box, above-right), as well as their corresponding app groups. Different desktops won't interfere with each other, but if you install several desktops and, in particular, their extra applications, you'll find your Applications menu becoming increasingly cluttered. And, most annoyingly, you'll have several identically named apps (the ones with generic names such as *Text Editor*). For a more pure experience, you might prefer to install instead the relevant Fedora spin (if available).

If you found the @games group, then you have several months' worth of distractions. If you were to install the whole thing (using `dnf group install --with-optional games`), then it would cost you a chunky 24GB. So, you probably shouldn't do that, but do check out some of the gaming offerings. All work and no play etc. You'll find *Steam* available in the *Software* app, too – you can install it from Flathub (though some people have issues here) or use the package from RPM Fusion. Both of these require you to have enabled third-party repositories (either from install, or from the settings in *Software*). We should also mention that *Lutris* is easy to install, too, which can manage your games from Steam, Epic Games, GOG, Humble and more. *Lutris* makes it easy to manage Proton versions

## » SPINNING OUT

If you really dislike Gnome, or just would rather have a different desktop, good news. There are several alternative spins of Fedora. Spins are official but not quite as thoroughly tested as the flagship Gnome release. The most popular spin is that of KDE Plasma, which in this incarnation uses Plasma 6.1 (6.2 has only just been released as we write this). KDE's Wayland support has come on in leaps and bounds, with Triple Buffering and Explicit Sync added in 6.1. And, yes, it does work with Nvidia cards. New in this outing is the addition of a Plasma Mobile spin. So, if you really want to see into the future, you can put Plasma on your phone. It also works on tablets and 2-in-1s.

Beyond the KDE ecosystem, there are also spins for several other desktops – (deep breath) Xfce, Cinnamon, MATE (with *Compiz* for fancy desktop effects), Budgie, LXDE (yes, it still lives), LXQt, Miracle (based on Mir), the minimal i3 tiling window manager, Sway (like i3 but Wayland), SOAS (Sugar on a Stick, the learning-friendly desktop from Sugar Labs), and Xfce.

But the leftfield editions don't stop there; besides desktop proclivities, there are also Labs, which cater to different use cases. You'll find Labs for Astronomy, Gaming, Security, Audio (in the form of JAM) and much more that you can enumerate for yourself at <https://fedoraproject.org/labs/>.

as well, so you can get any Windows games working like a dream.

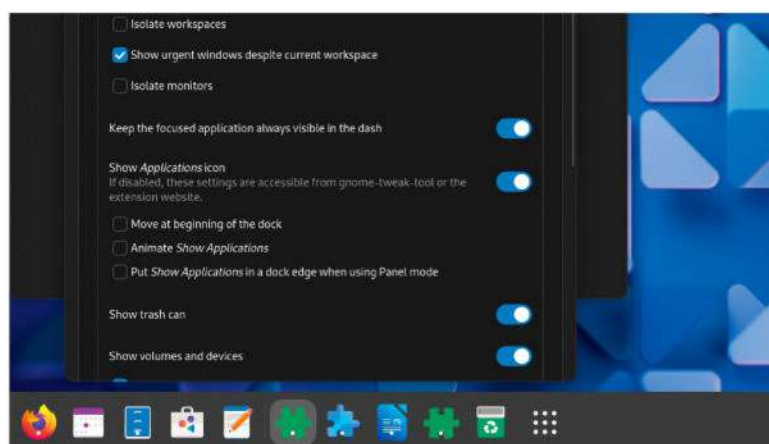
## The AI advantage

If you've got an Nvidia card and want to game, you'll want the proprietary driver. The open source Nouveau is excellent, and is also the only choice for older cards, but for AAA gaming, you want Nvidia's bits. The driver needs to build a kernel module, so you need to install compilers and all sorts. The following does it all:

```
$ sudo dnf install akmod-nvidia
```

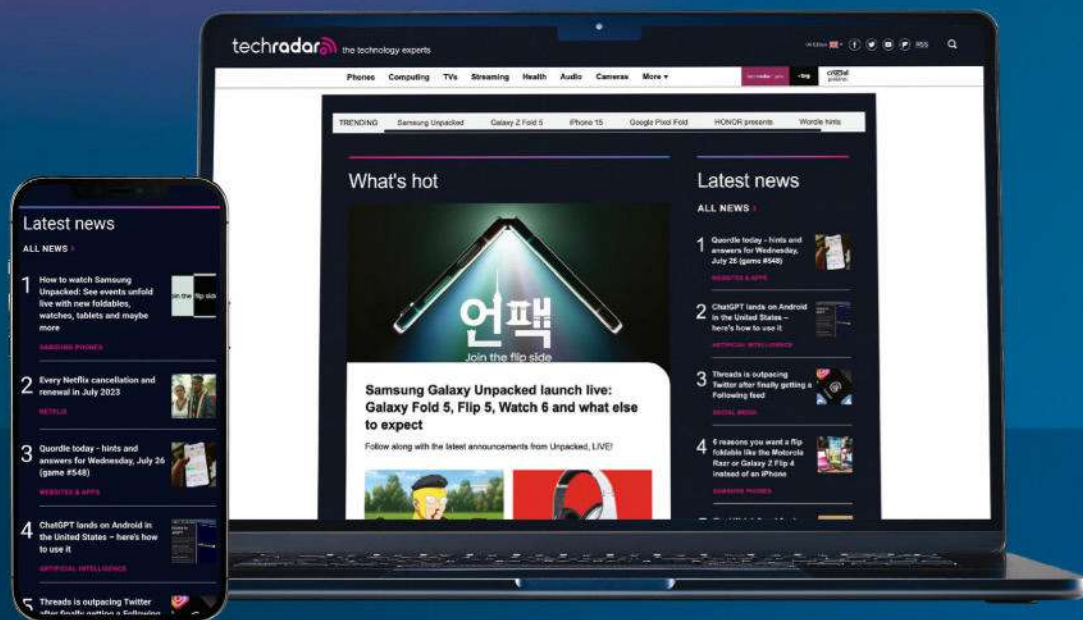
If you want to leverage AI and machine learning tools (following Fedora's marketing), you'll want **xorg-x11-drv-nvidia-cuda**. Depending on your purposes that might work on Wayland, too, but you may also have to switch to the X.org session (see previous page) in order to do deep learning with *PyTorch*, and so on.

We hope you enjoy Fedora. We've barely scratched the surface of this amazing distro, so let us know what you do with it. Or why you think it's better or worse than Ubuntu. We're off to tend to our gnomes. LXF



Dash to Dock's Panel mode sees the Dock extend to the edge of the screen. You can put the Apps menu at the edge, too.

# Meet the technology experts



- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

**techradar**  
the technology experts

**www.techradar.com**



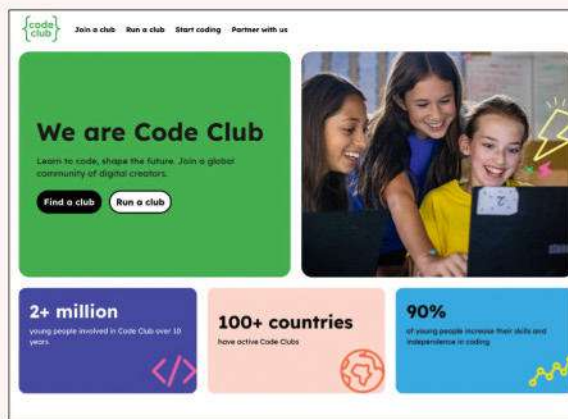
# Code Club and CoderDojo combine!

After years of success, the school clubs are combining into one single superpower!

**T**here's a new coding king in town and it's still Code Club! Since 2012, Code Club has helped two million young people get into coding and tackle technology in interesting ways. But nothing stays still, especially with technology, and the new vision is to boost Code Club's reach to 10 million young people worldwide over the next decade.

To help achieve this, a new more flexible model is needed, and that is now reflected in a new charter and new website at <https://codeclub.org>. The new system doesn't matter where you are, what you are, what community you're in, what schedule you need – heck, it doesn't even need to be called Code Club to be included! It's all about getting young people enjoying projects.

The highly successful CoderDojo will be folded into Code Club as part of the revamp, but if you're currently running a CoderDojo, nothing has to change. See [www.raspberrypi.org/blog/introducing-the-new-code-club/](http://www.raspberrypi.org/blog/introducing-the-new-code-club/).



Code Club goes from strength to strength.



**Les Pounder** works with groups such as the Raspberry Pi Foundation to help boost people's maker skills.

## » THE RADXA X4 IS HOT STUFF!

The Raspberry Pi has long been seen as the ideal for single-board computer (SBC) design. Just recently, however, I reviewed the Radxa X4, an Intel N100-based SBC that looks like a Raspberry Pi, but runs any X86 Linux distro or even Windows 11. Could it be a true contender for the Pi's throne? Well, yes, but it falls short.

The 4GB and 8GB models are priced to match the Raspberry Pi 5 4GB and 8GB (\$60 and \$80), but you need to spend a little more on the \$15 active cooler in order to keep the N100 from boiling itself alive. This is where the problem lies: the N100 is a potent four-core CPU that can turbo to 3.4GHz, but the Radxa X4 doesn't have the cooling power to allow the N100 to run at full speed. Don't get me wrong, it runs really well, but not to its full potential. I've tested another N100-based SBC, the LattePanda Mu, which uses a module-based design, like the Compute Module 4, and it manages much better performance.


The saving grace of the Radxa X4 is the GPIO, and this is where the Raspberry Pi Pico's RP2040 lends a hand. Built into the Radxa X4 is a Raspberry Pi Pico, with a full GPIO accessible over a serial connection. Linux works with the RP2040 with ease, Windows 11 took 90 minutes and a slew of drivers.

The Radxa X4 is a good machine, and well worth investigating if you need an X86 SBC in your next project. You can find out more at [www.tomshardware.com/raspberry-pi/radxa-x4-review/](http://www.tomshardware.com/raspberry-pi/radxa-x4-review/). [LXF](#)

## Bumper Pis

And faster SDs.

Out of the blue, the Pi Foundation released two new accessories: high-speed A2 SD cards and a rubber base bumper. The cards are made by Longsys, and were tested by running over 100,000 "surprise power cycles" while under a heavy I/O load, and offer exceptional random read and write throughput. <https://bit.ly/lxf322sd>



High-speed cards with bumper protection.

## Pi 500

Leaked GitHub post.

Over on the Raspberry Pi GitHub repository for the arm64 OS, we can see bcm2712-rpi-500.dts, which is the first device tree overlay for the Raspberry Pi 500. And it looks as though it will use the same Broadcom BCM 2712 SoC as the Raspberry Pi 5. Get full speculation from Tom's Hardware: <https://bit.ly/lxf322toms>



This is not a Pi 500!

# Anthias

Always one to draw attention, **Les Pounder** delves back into the world of digital signage and, of course, he is using his trusty Raspberry Pi.

## IN BRIEF

Previously known as Screenly OSE, this free Raspberry Pi-based signage OS is easy to use and powerful enough to turn any old screen into custom advertising signage. If you have a box full of old Raspberry Pis, they could be pressed into service for your next event.

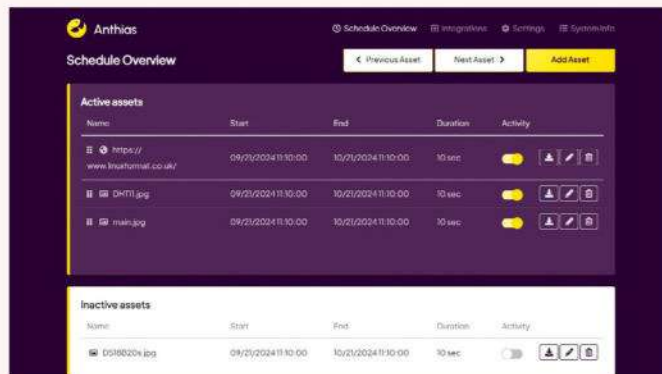
**F**ormerly known as Screenly OSE, Anthias is a digital signage OS for the Raspberry Pi and low-power x86 NUCs. The new name distinguishes the free version (Anthias) from the paid-for Screenly service. We've used Screenly OSE for conferences to deliver digital advertising and talk schedules across many Pis.

Unlike the older Screenly releases, Anthias is aimed at Raspberry Pi 2 and newer. You need the extra horsepower to play the myriad of media files. The team is working on fixing the installer for the Pi 1 and Zero, but we also can't see an installation candidate for the Pi 5. With that in mind, we tested on a Pi 4 8GB. Installation was via *Raspberry Pi Imager* but you can also install via a *Bash* script in Raspberry Pi OS Lite. This method offers the most flexibility, but you need to be comfortable with the terminal.

With the card flashed, we inserted it into the test Pi, powered up, then waited. It seems like Anthias needs to do some work behind the scenes, but we didn't know that, so we patiently waited for something to happen. When it did, we were instructed to go to another computer and open a browser to the IP address of the Pi 4. We connected our Pi 4 to the network via Ethernet; Wi-Fi is also an option, but you need to set this up before connecting the device.

The Anthias user interface is clear. The main interface is the Schedule Overview and at the top we have the active assets – these are the files or URLs that we want Anthias to display. At the bottom of the screen we have the inactive assets, which can be made active by clicking their toggle switch.

Adding an asset is easy. Click on Add Asset, then select a URL or an upload. Assets can be uploaded in bulk – just click on Add Asset, then Upload, then drag over the files for playback. Uploads can be images or video files. It is best to stick to MP4 video files, specifically those using H.264 encoding. We tested an



The default Anthias user interface is the Schedule Overview, and here we can administer the assets that will display on the digital signage.

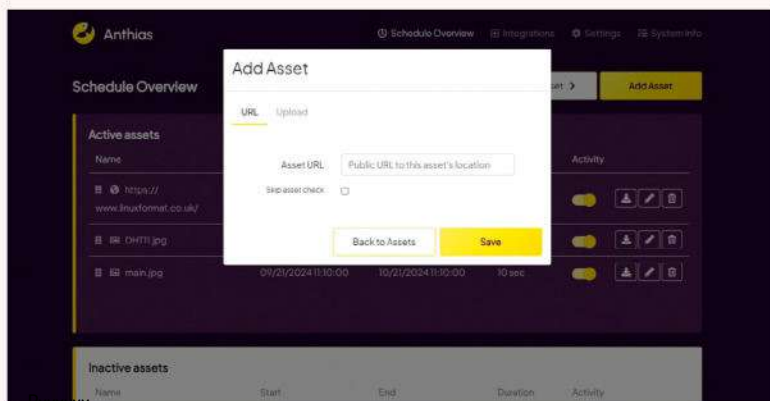
M4V file that didn't use H.264 and it failed to play. AVI files and other formats should be avoided as they will not work. This is where *HandBrake* comes in handy.

Active assets can be scheduled with a start and end date or time. Durations can be set, and the order in which assets are played can be tweaked by dragging them up and down the schedule. Adding a URL is a bit fiddly, but that's the fault of the underlying browser. Loading a heavy web page takes time, and depending on how it is constructed, there could be a slew of CSS, JavaScript and adverts that impact how it is displayed.

The Settings menu is where we get control of the underlying system. We can control audio output, date format and debug options. We can also back up the current setup, and restore that backup if needed. System Controls are the basic Reboot and Shutdown, nothing that could damage the system but they may be used by mischievous hackers, so lock down the install.

Anthias is polished and easy to use. Sure, it may be basic when compared to the paid-for service, but for offices and conferences, all we really need is a means to show images and basic web pages. Installation is a little clumsy – the long delay had us worried – but once it was ready, it was a great experience. **LXF**

Assets are broken down into images, videos and URLs. Images can be uploaded in bulk; URLs can be any web page.



## VERDICT

**DEVELOPER:** Screenly

**WEB:** <https://anthias.screenly.io>

**LICENCE:** GNU GPL V2

**FEATURES** 8/10

**PERFORMANCE** 9/10

**EASE OF USE** 9/10

**DOCUMENTATION** 8/10

Free and easy to use, Anthias is a great way to repurpose old Raspberry Pi for digital signage in schools and events.

» **Rating 9/10**



# Elegoo Mars 5 Ultra

When they said Mars needs women, this is exactly what **Denise Bertacchi** had in mind!

## SPECS

**Build:** 153x77x165mm  
**Screen** 7-inch monochrome, 9K (8,520x4,320)  
**Light:** COB light source  
**X/Y res:** 18x18 microns  
**Z-axis:** 0.02mm  
**Max speed:** 150mm/h  
**Exposure:** 2.2 seconds  
**Controls:** 4-inch touch panel  
**Comms:** USB, Wi-Fi  
**Extras:** Tilt release, auto-levelling, overheat, AI camera, warp detection, resin low  
**Size:** 260x268x451.5mm, 8.8kg

**I**f you're interested in 3D printing gaming miniatures, Elegoo's Mars 5 Ultra is precisely what the Dungeon Master ordered. You don't need to be a 3D printing expert to get started with this machine, thanks to its bevy of automatic features. It's also fast and precise, with an impressive 9K resolution.

The Mars 5 Ultra includes a tempered glass LCD screen protector – a huge benefit for those new to resin printing. It also boasts automatic levelling and has a tilt release feature that gently peels each layer off the vat's bottom film. It has a sensor to detect resin levels and failed prints to save you time and headaches.

Everything you need to get started is included, from scrapers to a printed manual and software on a USB stick. The printer comes fully assembled; you only need to remove the LED screen protector used for shipping, screw in the vat and build plate, and you're ready to print.

Unlike older resin printers, the Elegoo Mars 5 Ultra ships factory-levelled. Simply put the build plate in place and you're ready to go. It can do an automatic calibration before each print to make sure it's all in the right place.

The Mars 5 Ultra has a 9K resolution, which can be confusing because it refers to the total pixel count rather than a per inch or centimetre number. More meaningful is the 18x18µm XY resolution figure, which is how tiny the pixels are on its 7-inch LCD screen. A human hair averages 50µm, so we're talking about pixels less than half a hair. This gives us super-sharp details.

The build plate is laser-etched, which grabs on to prints a little too well. We had to dial back the first layer burn time in order to scrape prints off. The top of the plate has enough slope to allow most of the resin to drain back into the vat, but the screws are hidden under a bulky cover with standoffs that makes cleaning a chore.

The build surface is motorised and rocks the vat each time it prints a layer. This motion, combined with lifting the build plate, quickly and gently peels the model off the bottom film. The results are faster prints with a lower fail rate. It can also determine if there is enough resin to cover the bed and stops printing if it gets too low. The AI



**I** Compact and feature-full, the Elegoo is an ideal starter resin printer.

camera can also detect print fails, but only after the model is tall enough to be seen over the edge of the vat.

The Mars 5 Ultra ships with a copy of *ChituBox Basic* and is also compatible with *Lychee Slicer*. *ChituBox Basic* does everything you need, including hollowing the model and providing reliable automatic supports. The free version of *Lychee Slicer* is equally good, but includes ads, which can be annoying. Both have profiles for the Mars 5 Ultra, though the first layer needs a little less time.

Running *Rocket Bust* by Wekster, this printed in two hours 27 minutes, using a 2.2-second exposure time – 46 minutes faster than the Elegoo Saturn 4 Ultra and about the same time as the Anycubic Mono M7. The print is smooth and clean, with good details on the teeth and fur. This was printed using Elegoo's Rapid Standard Resin – you'll get better results with high-res resins.

The Elegoo Mars 5 Ultra is a great choice for gamers wanting to 3D print miniatures or beginners looking for a reliable first printer that will hold their hand. **LXF**

**I** Close-up of our Rocket test print.



CREDIT: Elegoo

## VERDICT

**DEVELOPER:** Elegoo  
**WEB:** <https://uk.elegoo.com>  
**PRICE:** £232

<b>FEATURES</b>	<b>9/10</b>	<b>EASE OF USE</b>	<b>9/10</b>
<b>PERFORMANCE</b>	<b>9/10</b>	<b>VALUE</b>	<b>9/10</b>

Elegoo knocks it out of the park with an affordable, easy-to-use resin 3D printer that's practically plug and play.

» **Rating 9/10**

## YOUTUBE

# Get live updates with the YouTube API

Seeking fame and fortune, **Les Pounder** wonders why he has only 586 subscribers on YouTube. Is it his Blackpool accent?



**OUR EXPERT**

**Les Pounder** is associate editor at Tom's Hardware and a freelance maker for hire. He blogs about his adventures and projects at <http://bigles.com>.

### YOU NEED

- > Pi Pico W
- > Breadboard
- > I2C LCD display
- > 4x F2M jumper wires
- > Code: <https://bit.ly/lxf322code>

**A**re you a budding YouTuber and want to keep an eye on your subscriber numbers? We've got just the thing. We're making a Pi Pico W-powered subscriber counter using a low-cost LCD screen. It updates every hour with the latest subscriber numbers. We're also learning about the YouTube API.

First log in to the Google Cloud Console (<https://console.cloud.google.com/welcome>), then click on the project selector drop-down. In the new window, click on New Project. In the next window, give the project a name, we chose **LXF Subs Project**. Now click Create. Next we need to enable API access for the project.

In the left-hand menu, click on APIs & Services > Library. In the search bar, search for YouTube Data API v3. When the result pops up, click on it, then on Enable. The last step is to obtain the credentials (API key), which enable our project to use the YouTube API.

Head back to the APIs & Services section and select Credentials. Click on Create Credentials. Choose API Key and in a few seconds an API key is generated. Do not share this key or put it in any version control (like GitHub). Do make a copy of the key in a text editor.

The final task is to get our channel ID. Open a browser window and go to YouTube; make sure you are logged in. Click on your face icon, top-right, and select Settings. In the next screen, click Advanced Settings. Copy the channel ID and keep it safe in the text file.

### Building the circuit

The LCD screen used in this project is fairly basic and resembles those used in vending machines. The I2C version of the HD4470 LCD display needs only four connections, much easier than the multiple pins in the non-I2C version. The circuit diagram in the download for this tutorial shows where to make the connections, and the table below provides a quick reference.

LCD Pin	Wire colour	Function	Pico pin
GND	Black	Ground reference	Any GND
VDD / VCC	Red	5V power	VBUS
SDA	Orange	I2C data	GP0
SCL	Yellow	I2C clock	GP1



The build is basic, but you could package this up into a 3D-printed case.

We assume you have already installed *Thonny* for the Pico. If not, Tom's Hardware has a guide: [www.tomshardware.com/how-to/raspberry-pi-pico-setup](http://www.tomshardware.com/how-to/raspberry-pi-pico-setup).

Open *Thonny* and connect the Pico to your machine. Go to Tools > Options > Interpreter. Set the interpreter to MicroPython (Raspberry Pi Pico) and set the Port to match the location of your Pico. Click OK. *Thonny* connects to the board and we now write code.

First, we need to install libraries (modules) to enable the Pico to talk to the I2C LCD screen.

We're using dhylands python\_lcd ([https://github.com/dhylands/python\\_lcd/](https://github.com/dhylands/python_lcd/)). Create a new blank file in *Thonny*, then open a browser to <https://bit.ly/lxf322api> and copy the contents to *Thonny*. Save the file to the Raspberry Pi Pico as **lcd\_api.py**.

Create another blank file and open another browser to <https://bit.ly/lxf322lcd>. Copy the contents to *Thonny*. Save the file as **pico\_i2c\_lcd.py** to the Pico.

Create another blank file to contain any passwords or API keys we need. In the file, create four objects for your Wi-Fi SSID, password, API key and channel ID. Replace the text inside the quotation marks with your Wi-Fi details – do not delete the quotation marks as this tells Python that the data inside are strings.

```
SSID = "YOUR WI-FI SSID"
PASSWORD = "YOUR WI-FI PASSWORD"
API_KEY = "YOUR YOUTUBE API KEY"
CHANNEL_ID = "YOUR YOUTUBE CHANNEL ID"
```

Save the file to the Raspberry Pi Pico as **secrets.py** and now we can move on to the main section of code.

Create a new file and start importing the modules of MicroPython code needed to make the project work. Network connects our Pico W to the internet, secrets is our file containing passwords and API keys. Time



controls the pace of the code, `urequests` and `ujson` are used to make web requests using the JSON format. Machine enables our code to talk to the GPIO and `pico_i2c_lcd` enables the use of the I2C LCD display.

```
import network
import secrets
import time
import urequests
import ujson
from machine import I2C, Pin
from pico_i2c_lcd import I2cLcd
```

The next bit of code handles setting up our screen. We need to tell the Pico where it is on the I2C bus. We are using the I2C channel 0 and we need to find the device on the bus. Each device has a unique ID so we'll scan the bus, get the address, then use it to connect. Then we tell the LCD module that we have a screen with two lines, 16 characters wide. Lastly we turn on a blinking cursor to confirm the connection is made.

```
i2c = I2C(0, sda=Pin(0), scl=Pin(1), freq=400000)
I2C_ADDR = i2c.scan()[0]
lcd = I2cLcd(i2c, I2C_ADDR, 2, 16)
```

Now we need to tell the Pico W that we want to get online, and what SSID and password it needs. The five-second sleep is optional, but we find it increases the chance of a good connection. Finally we print the connection status to the Python Shell.

```
wlan = network.WLAN(network.STA_IF)
wlan.active(True)
wlan.connect(secrets.SSID, secrets.PASSWORD)
time.sleep(5)
print(wlan.isconnected())
```


We create an object to store the URL of the YouTube API service.

```
YOUTUBE_API_URL = 'https://www.googleapis.com/
youtube/v3/channels'
```

Create the start of an exception handler process. This tries to run any code within it; any issues (exceptions) are handled later. The first part of the section is a **while True** loop to continually run the code.

```
try:
    while True:
```

Using an object, `url`, we need to create the custom URL that contains our YouTube API URL, channel ID and API key. We drop this information into the string using Python's F-strings formatting:



```
url = f'{YOUTUBE_API_URL}?part=statistics&id={secrets.CHANNEL_ID}&key={secrets.API_KEY2}'
```

Using this crafted URL, we use `urequests` to get the info from the YouTube API and store it in `response`:

```
response = urequests.get(url)
```

If we get a good response (HTTP 200), the info inside `response` is parsed using JSON formatting:

```
if response.status_code == 200:
    data = ujson.loads(response.text)
```

We want to extract just the subscriber numbers. As the data is stored using JSON, we need to drill down to the specific value using a series of keys:

```
subscriber_count = data['items'][0]['statistics']
['subscriberCount']
```

We print the subscriber count to the Python shell:

```
print("Subscriber count: " + subscriber_count)
```

To display the information on the LCD screen, we first need to turn on the backlight, then put a string of text on the top line: **"Subscribers"**. The `\n` is a Python escape character to start a new line after this string:

```
print("Subscriber count: " + subscriber_count)
lcd.backlight_on()
lcd.putstr("Subscribers\n")
lcd.putstr(subscriber_count)
```

If there is a response other than HTTP 200, the **else** part of the code activates, printing a message to the Python shell, before waiting one hour (3,600 seconds), clearing the LCD screen, then the main loop repeats:

```
else:
    print('Error fetching data from YouTube API')
    time.sleep(3600)
    lcd.clear()
```

The final part of the code is the exception handler:

```
except Exception as e:
    print(f'An error occurred: {e}')
```

Save the code as **main.py** to the Pi Pico. In *Thonny*, click on the run button to test the code. **LXF**

Just four wires connect the large LCD display to the Raspberry Pi Pico W via an I2C interface.

## » GET THE PICO W ONLINE – GRACEFULLY

In the project code, we were a little basic in how we connected to Wi-Fi with the Pi Pico W, so let's make something a little better.

We start by creating the `wlan` object and activating the Wi-Fi:

```
wlan = network.
WLAN(network.STA_IF)
wlan.active(True)
```

If there is no connection, the response is **False**. The Pico W tries to connect again:

```
while wlan.isconnected() == False:
```

```
    wlan.connect(SSID, PWD)
```

It prints the connection status to the Python shell and to the LCD display, before sleeping for two seconds, then clearing the LCD:

```
print("No connection")
lcd.putstr("No
connection")
sleep(2)
lcd.clear()
```

If the connection is made, the **else** condition activates, creating an object, `ifconfig`, to get connection details:

```
else:
    ifconfig = wlan.ifconfig()
```

We turn the LCD backlight on and print the IP address to the Python shell and display, before sleeping for five seconds and clearing the LCD.

```
lcd.backlight_on()
print(ifconfig[0])
lcd.putstr("IP ADDRESS")
lcd.putstr("\n"+ifconfig[0])
sleep(5)
lcd.clear()
```

» GET YOUR Pi FILLING HERE Subscribe now at <http://bit.ly/LinuxFormat>

## PHOTOGRAPHY

# Use a Pi to capture infrared photos

**Mile Bedford** unveils an unusual but attractive form of photography courtesy of a Pi NoIR camera and suitable post-processing.



### OUR EXPERT


**Mike Bedford** has dabbled with infrared photography for some time. As he's discovered, though, using a Raspberry Pi can open up some exciting new possibilities without breaking the bank.

**I**nfrared is commonly thought of as heat, but that's only true of far infrared. Closer to visible light – in fact, immediately adjacent to the red end of the visible spectrum – is so-called near infrared. And while far infrared can only be photographed using special cameras, near infrared can be photographed with ordinary cameras, more or less. So, near infrared photography, which we will refer to as simply infrared photography, is much more accessible.

You could use most phones or ordinary cameras for infrared photography (see **LXF248**) but here we're using Raspberry Pi NoIR cameras. For reasons we'll see, these offer several benefits, including the ability to create attractive false colour infrared images. After providing a quick overview of infrared itself and infrared photography, we'll move on to a couple of practical themes. First we'll see how to take infrared photos, and then we'll look at the crucial subject of post-processing. So, if you find the prospect of photographing the invisible intriguing, and are prepared to spend a very modest amount of money, you have come to the right place.

### Infrared – the basics

It's good to know something about infrared, and not just out of a desire for knowledge. Infrared, as well as visible light, are both a type of electromagnetic radiation and they are differentiated by their wavelengths. The visible spectrum ranges from purple light, which has a wavelength from around 400nm, through indigo, blue, green, yellow and orange, to red light, which has a wavelength up to around 750nm. Immediately above this we find infrared, which is considered to cover a huge range of wavelengths, from 750nm to 1mm. Of this, near infrared covers wavelengths of 750nm to 1,400nm.



Don't underestimate the appeal of black and white photography, and infrared makes it even more spectacular.

Although our eyes can't see infrared, the imaging sensors in our cameras and phones are sensitive to it. But it's not quite that simple. The infrared sensitivity of camera sensors would cause ordinary photos to have an unnatural colour cast, so manufacturers place a filter over the sensor to prevent infrared from reaching it. Normally it's not a total show-stopper, though, as the filter isn't 100% effective. And it's possible, with many cameras, to take advantage of that to take infrared photos. However, to do so it's necessary to prevent visible light reaching the sensor, and this requires the use of a filter over the lens – more on that later.

Using an ordinary camera in this way is a cheap and easy way of starting out in infrared photography and we looked at this method in **LXF248**. However, there are a couple of drawbacks. First, because of the internal infrared blocking filter, very little infrared reaches the sensor. This means that exposure times are long – running to several seconds – even on a bright sunny day. Accordingly, handheld photography just isn't an option, so you have to use a tripod. Second, that internal filter, together with the filter over the lens, severely limits the range of wavelengths that are recorded. While this doesn't affect classic black and

### QUICK TIP

Reference to wavelengths might seem positively archaic, with radio waves having been referred to by their frequencies instead of wavelengths for several decades. Wavelengths are still commonly used when referring to light and infrared, though.





A 720nm filter and false colour processing keep the drama of black and white infrared while producing blue sky and subtle colouring to the trees.



With sufficient ingenuity in post-processing, photos taken with a 720nm filter can exhibit rather more saturated colours.

white infrared photography, it does prevent you from trying your hand at false colour infrared photography.

### The Raspberry Pi option

Our aim is to introduce you to an alternative way of carrying out infrared photography that overcomes the obstacles we've just seen. This involves using a camera that doesn't have an internal infrared blocking filter, and is the method of choice for serious infrared enthusiasts. However, this is expensive because such cameras are very rare in the extreme, so it's necessary to modify a camera by removing the filter. In turn, this is expensive, as it's not an easy job, so you'd be advised to have the job done professionally, and it renders the camera unsuitable for ordinary photography.

However, there's one notable exception, namely the family of Raspberry Pi cameras. If we exclude, for the moment, the HQ Camera, all generations of the Pi Camera have been available with the option of a NoIR variant. Those letters indicate that these models have no infrared (IR) blocking filter. In the main, we suspect, these cameras are used for night-time photography in conjunction with infrared LED lighting, but they're also suitable for artistic infrared photography. And although the HQ Camera isn't available in a NoIR version, it's not really very difficult to remove its filter, and instructions on how to do that modification are available on the Pi

website ([www.raspberrypi.com/documentation/accessories/camera.html](http://www.raspberrypi.com/documentation/accessories/camera.html)). Do be aware, though, that this invalidates the warranty.

We previously referred to the need to use a filter over the lens when taking infrared photos, but we glossed over it, so here's the low-down. Our aim is to capture a scene using only infrared, not a combination of infrared and visible light, the latter of which would look like a colour photo with somewhat unnatural colours. To do this, we need to use an infrared filter over the lens. However, unlike the internal infrared filter, which blocks infrared, we need a filter that blocks visible light, allowing only infrared to pass. These filters are available in several variants – differentiated by their cut-off wavelength – and we discuss these in the boxout (page 49). Filters are fixed in front of a lens by screwing them into the filter thread that is present on most cameras (but not phones), and this enables you to swap between different filters. However, with the exception of the HQ Camera lenses, the lenses on Pi cameras don't have filter threads. We'll look at options for overcoming this drawback later as part of our discussion on how to build a Pi infrared camera.

### Build an infrared camera

Because a portable Pi-based infrared camera will be almost identical to an equivalent visible light camera,

#### QUICK TIP

Don't fall into the trap of thinking of black and white as the poor relation to colour photography. Black and white is an art form in its own right and can produce spectacular results in cases where a colour photo wouldn't have nearly the same appeal.

## » THERMAL INFRARED

Thermal infrared, aka far infrared, doesn't have a lot in common with near infrared, the subject of this tutorial. What it does have in common, though, is that you can photograph it with a Pi. If you go down that route, it'll be for scientific applications or just to learn about thermal imaging – not really as an artistic pursuit – but it's no less interesting for that.


Because thermal infrared is much further from visible light than near infrared, we can think of it as heat, so although we can't see it, we can feel it on our skin. A thermal infrared camera can better detect direction than our perception of heat, so a far infrared photo can show temperatures across a scene. These are normally presented as false colour with black,

blue, green, yellow, orange and red representing increasing temperatures.

The downside of thermal infrared photography is its resolution, with 640x480 (0.3MP) being typical, even for professional cameras. However, if you already have a Raspberry Pi, an add-on camera need not cost a lot. The Pimoroni MLX90640 Thermal Camera Breakout

costs around £60. You're not going to be creating high resolution images with its 32x24-pixel sensor, but it's great for experimentation. And if you also capture an ordinary monochrome image with a standard Raspberry Pi camera – bearing in mind the difficulties of parallax – you can combine them to paint the IR colour on a higher resolution image.





dramatic results require an 850nm filter (*see box on opposite page*). And if they're taken of a suitable subject, they can, indeed, be very dramatic. To get the maximum dramatic effect, though, you need to take your photos in bright sunlight. The element of drama comes from the fact that infrared black and white photos exhibit very high contrast. In particular, blue skies often appear very dark – almost black at times – while grass and the foliage on trees come out a vibrant white. That white could be described as almost glowing, and this gives rise to an eerie, yet strangely attractive appearance. This fact will probably also influence which seasons of the year to concentrate on. Certainly experiment with other subjects

Straight out of the camera, a photo taken with a 720nm filter looks rather uninspiring. Post-processing can change all that, though.

we're not going to give you detailed instructions on how to build such a camera. We will give you some guidance, but the space saved will allow us to give our discussion on the use and post-processing of infrared photos the attention it deserves.

The bottom line is that Raspberry Pi camera projects are by no means in short supply, so you'll easily find them online. In your searching, it's probably a good idea to include "portable" or "handheld" as a search term, otherwise you'll find loads that don't meet your needs. These designs generally include a Raspberry Pi board, a Pi camera – which is generally a V1, V2 or V3, although HQ Camera designs also exist – an LCD display, a battery, a case with a shutter release button, and the software. In fact, if you fancy hacking your own code, or modifying the code supplied with a published project, it's a fairly simple exercise, thanks to the availability of camera libraries.

## Behind the lens

Taking a photo with your Pi-based infrared camera is much the same as with the visual light alternative. What differs, though, is the effect that you'll get and this will influence your choice of what to photograph and when to take that photo.

We'll look, first of all, at the traditional black and white infrared photo. You can take these photos with just about any visible light blocking filter, but the most

at other times of the year, but if you really want to maximise that glowing white look, you need to pick seasons when there are leaves on the trees.

If you're taking a photograph with the intention of post-processing it to achieve a false colour effect, you'll probably choose a 720nm filter (*again, see the box on filters, opposite*). However, if you want to leave the option of black and white or false colour open, you can still shoot with a 720nm filter. Indeed, because a 720nm filter enables you to work with black and white or false colour, if you're only going to buy one filter, that should be your choice. The black and white results might not be quite as high in contrast as with an 850nm filter, but they can still be quite spectacular. Much of what we said about black and white photos still applies to false colour. So, skies will be dark, but probably dark blue as opposed to black, and greenery will be very light, but might be a very pale yellow or orange rather than pure white. And this, of course, will continue to impact the times of day and seasons of the year you choose.

## Post-processing

Post-processing is an essential part of infrared photography, as indeed it is if you want the best results from any form of photography. We used the *GIMP* photo-editing package, although you'll be able to do the same with pretty much any fully-featured software. And while it's a must for false colour work or for shooting black and white with a 720nm filter, don't ignore it for all black and white photos. After all, this can make the difference between an OK photo and a really good one.

Here we're starting with black and white. With an 850nm filter, photos straight from the camera will appear largely monochrome. If you're using a 720nm filter, though, your photos will have a reddish tint to them, so it's absolutely essential to convert the photo to black and white. Even if you shot with a 850nm filter, though, it's better to do this because a genuinely black and white photo will look decidedly better than an approximately black and white one. There are several ways to do that conversion, the easiest being to

### QUICK TIP

It's perhaps questionable whether a Raspberry Pi warrants it, but most infrared photography aficionados would argue that you really ought to work with RAW images instead.

Your Pi-based infrared camera won't look much different from any other Pi handheld camera except for the filter on the front. Internally, though, you'll be using a NoIR variant of camera.





change the mode to greyscale, if available, or else desaturate it. However, because there's more colour information in photos taken with a 720nm, this gives you more options, so you might like to experiment. Bear in mind, though, that not all photos benefit to the same extent from a particular approach. One approach is to select what proportion of the three primary colours to use in generating the greyscale image. In *GIMP* you can find this at Colours > Desaturate > Mono Mixer, and an obvious first thing to try is using just each of the primary colours in turn.

Making sure you've got a properly monochrome image is just the start and you can probably further improve your photo. And we'd suggest that one option is by far the most important one to try. You might choose to do more, and the sky's the limit, but do try this first. We described the infrared effect as being characterised by high contrast, and so it is. However, that doesn't mean you can't give it a bit of extra help, just so long as you don't overdo things. Most photo-editing tools enable you to increase the contrast but you'll get far more control using the Curves tool, which you'll find in *GIMP* in the Colours menu. This maps input intensities to output intensities and often does wonders for any form of photography. The classic approach, and the one you should try first, is to lighten the bright parts of the image and darken the darker parts, while leaving most of the intermediate values untouched. Selecting an S-shaped curve is the way to do this – read up on general photo processing if you're unfamiliar with this.

## Fake news

Next we come to false colour, which involves converting the rather uninspiring reddish image as it comes out of the camera with a 720nm filter (or a filter with an even lower cut-off wavelength) to something more appealing. First of all, because the unprocessed image will have little if any true white, we need to correct the colour cast. The easiest way to do that in *GIMP* is to go to Colours > Levels, and then, in the Adjust Colour Levels dialog box, select the white point picker that appears below All Channels. Now just click on the image on a place that should appear white, most probably a bright area of grassland. The excessive



colour cast should now have been reduced, but there'll still be an abundance of rather unpleasant red or purple colouration, especially in what you'd expect to be a blue sky. To correct this, we need to swap the red and blue channels, which we can do via Colours > Components > Channel Mixer. In the dialog box, change Red in the Red channel to 0.000 and Blue in the Red channel to 1.000. Similarly, change Red in the Blue channel to 1.000 and Blue in the Blue channel to 0.000. And, as if by magic, the sky will have turned blue and, overall, the appearance of the scene will be much more aesthetically pleasing.

This is the general workflow for the initial processing of a false colour infrared photo, and the overall appearance will differ depending on whether you used a 720nm filter or something with a lower cut-off. With a 720nm filter, though, you might just find that trees are faintly purple and you might simply prefer the orangey yellow that is common with lower cut-off wavelength filters. If so, experiment with altering the hue of the magenta tones at Colours > Hue Saturation. This effect is fairly subtle with a 720nm filter, but that's not necessarily a bad thing. However, using a lower wavelength cut-off filter results in more saturated yellows or oranges. You might also want to try the Curves tool that we looked at for processing black and white photos, but don't go too far. There are lots of other options to try, so experiment away in the search for that perfect infrared image. **LXF**

Unlike near infrared, far infrared is effectively heat, so a thermal IR camera can produce a temperature image.

## » INFRARED FILTER OPTIONS

In looking at infrared filters, we're thinking about the type you'll put over the lens and that allows infrared to pass while blocking most visible light. These are defined by their cut-off frequency, so an 850nm filter blocks anything below 850nm, thereby cutting out all visible light and even the bottom end of near infrared. This is an ideal filter

for black and white work but it won't work for false colour.

Next we come to the 720nm filter. At first sight it appears totally black, but if you hold it up against a bright light, you'll observe that light as a faint red image. This is because it allows the very top end of the visible spectrum to pass through, plus all the near infrared. This type of filter

works for black and white photography, although the effect isn't quite as dramatic as with an 850nm filter. It's also good for false colour, although the colours of foliage and grass can be subtle, but perhaps no less attractive for that. For this reason, it's a good all-rounder.

We're not done yet, though; moving further down the

spectrum, we come to filters with a cut-off around 650nm or 590nm. These actually look red because they allow more of the visible spectrum to pass as well as all the near infrared spectrum. These are sometimes referred to as enhanced false colour filters or similar, and that's just what they do. They result in more saturated false colour images.

# MEDITATING ON ZEN 5

**T**he AMD Zen 5 CPU architecture will form the backbone of the next generation of Ryzen desktop and laptop processors, and while we got a general idea of the new tech during Computex earlier in 2024, we've now been treated to the full architecture monty.

Before diving into the finer details, it's worth revisiting the information revealed back at the start of June. At Computex 2024, AMD announced the Zen 5-based Ryzen 9000-series desktop CPUs and Ryzen AI 300 mobile APUs were set to launch during the month of July. As was expected, Dr Lisa Su's Computex keynote was full of general details for a general audience, with little meat on the bone regarding the architecture itself.

The Ryzen 9000-series is code-named Granite Ridge. And, apart from the architectural improvements, the new models are physically similar to their Zen 4 (AM4) Ryzen 7000-series predecessors.

The CCDs, or Core Complex Dies, are still made up of eight cores, with a shared 32MB of L3 cache. The Ryzen 9 9950X and 9900X come with two CCDs, while the Ryzen 7 9700X and Ryzen 5 9600X only

need one. All chips include a separate I/O die, which is the same as that used with Zen 4.

Granite Ridge includes an RDNA2 integrated GPU with two compute units. The CCDs are manufactured with TSMC's N4 process, while the I/O die is fabbed on TSMC's N6 process.

Strix Point mobile APUs come with monolithic dies, and it's a large one at that. They're being branded as the Ryzen AI 300 series. You just know that AI had to be placed in that name, didn't you? Two models have been announced so far. They are the Ryzen AI 9 HX 370 and Ryzen AI 9 365. These feature 12 cores/24 threads and 10 cores/20 threads respectively. Both chips' CPU cores are based on the Zen 5 architecture. Both include an XDNA 2 NPU capable of up to 50 TOPS, and RDNA 3.5 integrated graphics.

The Ryzen AI 9 HX 370 features Radeon 890M graphics with 16 CUs, while the Ryzen AI 9 365 features Radeon 880M graphics with 12 CUs. AMD demonstrated laptops with 890M graphics running the very demanding *Cyberpunk 2077* at 1080p at above 55fps, albeit with FSR and AFMF activated.

As they are Ryzen 9 models, these chips will find their way into high-performance notebooks. Ryzen 7 and Ryzen 5 models are sure to follow later in the year.


Strix Point features a mix of higher performing Zen 5 cores and efficiency focused Zen 5c cores. These aren't like Intel's hybrid P and E cores, though; Zen 5c cores are the same as the larger cores, though they come with less L3 cache and lower clocks. This saves die space, allowing AMD to allocate its transistor budget to other areas, such as the NPU and GPU. AMD says this approach doesn't require software awareness or anything like an on-chip Thread Director that Intel

## AMD Ryzen 9000-series

CPU model	Arch	Cores/ threads	Boost clock	L2+L3 cache	Default TDP
Ryzen 9 9950X	Zen 5	16/32	5.7GHz	80MB	170W
Ryzen 9 9900X	Zen 5	12/24	5.6GHz	76MB	120W
Ryzen 7 9700X	Zen 5	8/16	5.5GHz	40MB	65W
Ryzen 5 9600X	Zen 5	6/12	5.4GHz	38MB	65W



Discover what the new AMD architecture can do for you and your Linux kernel with Chris Szewczyk.



## » MOTHERBOARD CHIPSETS

AMD has moved forward from its 600-series chipsets for Ryzen 7000 to the 800-series for Ryzen 9000, but the new chipsets still leverage the AM5 LGA1718 socket. The 800-series motherboard models are arriving in 2024 into 2025. The Ryzen 9000 chips are fully compatible with existing AM5 motherboards (after the requisite BIOS update), so there are plenty of options for early adopters. The staggered launch is obviously not ideal, but at least support for future chips will be around for the long haul – AMD plans to support the AM5 socket until 2027-plus.

The upper-tier X870 and X870E chipsets come with welcome improvements, but those advances stem from new mandatory feature requirements instead of expanded chip connectivity. For instance, the PCIe 5.0 interface is now standard on

the X870 lineup for both storage and graphics, whereas it was previously limited to the E-series boards. All X870 boards will now also have the USB 4.0 60Gb/s interfaces courtesy of third-party controllers, such as the ASMedia ASM4242 (the controller will consume some of the PCIe lanes from the CPU to operate at full bandwidth). The 800-series is built around the same Promontory 21 chipset silicon from ASMedia as the 600-series.

The high-end boards are nice, but the Ryzen 5 9600X is a natural pairing for a competent B850 motherboard. AMD's B850 series features largely the same feature set as the existing B650 motherboards, with full support for overclocking both the processor and memory, and support for USB 3.2 20Gb/s. AMD hasn't yet revealed the

full breakdown of the ports, but aside from optional USB 4 support, we've gleaned that the B-series feature set remains the same.

In addition to the lower-tier B850 boards we expected, AMD has added a new tier for the B-series with the new B840 series. This chipset bridges the gap between the A and B-series to help reduce platform costs. The B840 series supports memory overclocking, but CPU overclocking isn't

supported. They will also come with only the PCIe 3.0 interface, a big reduction from the mix of PCIe 5.0 and 4.0 interfaces on the B850 family. The B850 and B840 also get another cost reduction via support for USB 3.2 instead of the mandatory USB 4 on the X-series motherboards. We don't know how much the reduced feature set will impact pricing, but the B840 lineup seems best suited for OEMs and office machines.



A wide range of Zen 5-ready mobos are now out, with budget options to come.

## The AMD 800 Series Chipset Family

### Simplifying the value proposition for users

	PCIe <sup>®</sup>	USB	Overclocking	Graphics	Competition
AMD X870E Chipset	Gen 5 Graphics and NVMe	USB 4 Memory	CPU and Memory	1x16, 2x8	Z790
AMD X870 Chipset	Gen 5 Graphics and NVMe	USB 4 Memory	CPU and Memory	1x16, 2x8	X670
AMD B850 Chipset	Gen 5 NVMe (Optional) Gen 4 Graphics	USB 3.2 20 Gbps	CPU and Memory	1x16, 2x8	B760
AMD B840 Chipset	Gen 3	USB 3.2 10 Gbps	Memory only	1x16	B760

We all love more bandwidth, and the load and store functions of CPUs are no different. The 12-way L1 data cache is now 48Kb in size, up from the eight-way 32Kb of Zen 4. That's a whopping 50% increase, and AMD engineers did it without any latency increase, which is usually the case with larger caches. This L1D cache supports four loads per cycle; AMD says this is particularly beneficial for 512-bit loads. The L1 to L2 bandwidth has also doubled and the prefetching algorithms have been tweaked.

The final key Zen 5 improvements are to the floating point/vector math units. Zen 5 now supports native AVX-512. Zen 4

achieved this by 256-bit double pumping. Interestingly, AMD says Zen 5 processors will run AVX-512 without any frequency penalty. Strix Point supports full AVX-512, too, which will remain challenging for notebook cooling and TDPs, though AMD says it still supports the 2x 256-bit option depending on the configuration. Note that several of the particularly favourable pre-release AMD benchmarks take advantage of full AVX-512. Machine learning is another field that will benefit from AVX-512 instructions.

So, the Zen 5's front-end has been dramatically improved, as have the execution units. Latency is lower, while throughput and bandwidth have increased, all while keeping power consumption in check. In the case of Granite Ridge, all the SKUs have lower TDPs than their predecessors, with the exception of the Ryzen 9 9950X, which is the same. AMD is understandably quite proud of its 16% average IPC uplift. The gaming improvements are likely to be less overall, however, but do note the *League Of Legends* uplift at an impressive 21%.

The lower TDPs give users some additional Precision Boost Overdrive headroom, particularly in the case of the 9700X, where AMD claims users can gain as much as 15% higher performance. You'll get even more if you use *Curve Optimizer* and the new *Curve Shaper* tool. The latter allows users to tweak the underlying voltage curves to maximize performance.

### No NPUs

One of the interesting omissions from Ryzen 9000-series processors is a dedicated NPU. Why, you might ask? We asked AMD corporate fellow and chief Zen architect Mike Clark about this, and the answer was illuminating. Including an NPU wasn't deemed essential given that most Ryzen 9000 systems will include discrete GPUs, which are already highly capable AI processors. Granite Ridge also includes the AVX512-VNNI (Vector Neural Network Instructions) instruction set. AMD believes a complete redesign of the I/O die was therefore not justified. Time will tell if this design choice was the right one.

Zen 4 employed a dual-issue AVX-512 pipeline ("double-pumped" in AMD lingo), meaning it issued an AVX-256 instruction twice across a 256-bit interface to provide most of the performance benefits of AVX-

The budget B840 will be targeting OEM box shifters.

designs do. The 5c cores simply kick in lower on the voltage/frequency curve, when the larger cores would be power limited anyway.

### AMD Zen 5 architecture

AMD attributes most of the Zen 5 improvements to four key areas. They are instruction fetch and decode, integer execution, load and store, and the floating point/maths execution units.

The front-end redesign includes improved branch prediction with lower latency, better accuracy and greater throughput. Branch prediction is one of the fundamental pillars of x86 performance. The better it is, the fewer clock cycles get wasted, leading to better performance and power efficiency. The aim is "to keep the beast fed", as AMD colourfully put it. Downstream,

## TO DIE FOR?

"AMD believes a complete redesign of the I/O die was therefore not justified. Time will tell if this design choice was the right one."

Zen 5 includes dual port instruction and op caches, so it includes not just better branch prediction, but more predictions per cycle. Zen 5 also includes dual decode pipes, whereas Zen 4 included only one.

As improved predictions and instructions are coming through the wider pipeline, AMD needed to make improvements to its instruction dispatch and execution engine. This engine now includes eight-wide instruction dispatch and retire capability each cycle, an increase from the six of Zen 4. With more instructions comes the need for improved scheduling, and to this end AMD redesigned its ALU scheduler, which is now more unified than that of Zen 4. Zen 5's retire queue/reorder buffer is 40% larger than the 320 ops of Zen 4 at 448 instructions deep, giving the CPU a wider window of instructions for out-of-order execution. The Arithmetic Logic Unit (ALU) count has been increased to six from Zen 4's four.



512 while avoiding large impacts to die area and frequency drops.

Zen 5 supports the full 512-bit data path to deliver doubled AVX-512 and VNNI throughput, but AMD says the chip still runs at the same frequency as integer work for any given multicore workload during AVX-512, a by-product of its careful balancing act with the power characteristics of the integer instruction path. That stands in stark contrast to Intel's AVX implementation, which results in frequency reductions.

Other general improvements include a dual-pipe fetch and better branch prediction accuracy, dual-ported instruction and operation caches, a dual four-wide decode path, and an eight-wide dispatch. Zen 4 had a 32KB data cache, which has now been expanded to a 48KB 12-way L1 data cache (L1D) for Zen 5.

Ryzen 9000 series chips reportedly have a 15% thermal resistance improvement, leading to an average 7°C reduction at the same TDP. It's no secret that Ryzen 7000-series chips got a bit toasty, so this is a welcome improvement and it should allow 9000-series chips to hold higher clocks without thermal throttling with less-than-stellar cooling solutions.

One of the more eyebrow raising features is support for on-the-fly memory overclocking. For example, you could increase your memory speed for gaming, while lowering it for daily use.

## Zen 5 overclocking

DDR5 support has increased to 5,600MT/s officially, with EXPO profiles up to 8,000MT/s also supported. If you recall at Computex, G.Skill demonstrated a Ryzen 5 8500G system running DDR5-10600 memory. AMD says Granite Ridge should be capable of even more. We heard DDR5-11000 mentioned, though that is likely out of reach without extreme cooling. AMD says the sweet spot for Zen 5 remains at 6,000-6,400MT/s in terms of compatibility, cost and performance.

But, as DDR5 speeds are continually on the rise, it will be interesting to see whether memory in the 8,000MT/s-plus range could make sense for gaming, as higher frequencies will eventually deliver enough



raw speed to overcome the latency penalty you get when running 1:2 mode on the memory controller.

## Linux support

The good news is that the Zen 5 should offer excellent Linux support. Older AM5 motherboards may require a BIOS update to tweak any minor firmware issues, which is something you should check for when choosing a new system, either pre- or self-built. But largely support is ready to go in the kernel down to the mobile Strix Point GPU, as it's using the older RDNA2 graphics that have been open source for years now.

Not an issue for most home users, but Zen 5 compiler support was only added to GCC 14.1 back in April 2024. As Ubuntu 24.04 ships with GCC 13, hitting its release would have been far better, but at least it's publicly available. It wasn't even released for LLVM/Clang until mid-September, which raised eyebrows, though it was due to some legal hurdle rather than anything technical. But ultimately, support is now there with LLVM 19 and 20. That's enough theory, let's see how they perform over the page!

The main Core Complex that contains the CPU cores sits on a separate die from that of the I/O die.


## » FUTURE OF ZEN

We asked AMD about how the Ryzen 7 7800X3D compares to the Ryzen 7 9700X. The answer given was that the 9700X should squeak ahead on average, though cache-loving games will favour the 7800X3D.

On a related note, AMD pointed out that the 9700X is on average 13% faster than the Ryzen 7 5800X3D. This will give AM4 gamers a little incentive to make the switch, though 9000X3D chips will be worth waiting for if you are on the fence about switching to AM5. Of

course, AMD gave an entirely expected "no comment" when asked about 9000X3D chips, but you can be sure such chips are coming. AMD will be hoping to take some wind out of the sails of Intel and its upcoming Arrow Lake processors, which.

We picked up that AMD designed Zen 5 to be a foundational architecture. The front-end improvements were designed to do what AMD called "keeping the beast fed". Though AMD can't be expected to comment, we wonder



The Ryzen 7 7800X3D (left) still has the edge over the the Ryzen 7 9700X (right) for cache-loving games.

whether some of the groundwork is there to enable Zen 6 to scale to higher core counts.

Only AMD will know what bottlenecks can be improved upon, or if there is some low

hanging fruit waiting to be picked. The latter is unlikely given the historical difficulty of achieving significant inter-generational x86 IPC improvements along with the slowdown of Moore's Law.

# AMD Ryzen 5 9600X

Let's face it, this is the Ryzen you're going to buy.

## SPECS

**Socket:** AM5  
**Arch:** Zen 5  
**Process:** 4nm TSMC CPU, 6nm I/O  
**Cores:** 6  
**Threads:** 12  
**Clock:** 3.9GHz (5.4GHz boost)  
**Cache:** 480KB L1, 6MB L2, 32MB L3  
**Unlocked:** Yes  
**GPU:** AMD Radeon 2 Cores  
**Mem:** DDR5-5600, ECC mobo support, two-channel, 192GB max  
**PCIe:** v5 28-lanes  
**TDP:** 65W

**T**he first salvo of AMD's Zen 5 Ryzen 9000 processors are here. The new chips span from the £269 six-core 12-thread Ryzen 5 9600X to the £594 16-core 32-thread Ryzen 9 9950X, covering the same segments as the prior-gen Zen 4 Ryzen 7000s. The chips have the same core counts as their predecessors, with a new architecture offering significant performance and efficiency improvements. The Ryzen 9000 chips drop into the existing AM5 socket, which AMD will support until at least 2027.

The Ryzen 5 9600X replaces the Ryzen 5 7600X in the £300 price bracket. Pricing-wise, the 9600X contends directly with Intel's Core i5-14600K, a lacklustre Raptor Lake Refresh family member that offers little over its predecessor, the Core i5-13600K.

The Ryzen 5 9600X gets a slight 100MHz clock rate increase to a 5.4GHz boost, but the company also dialled the base clock back by 800MHz, which helps reduce the TDP. The move to the TSMC 4nm process and the efficiency of the Zen 5 architecture allowed AMD to drastically reduce the 9600X's rating to 65/88W, a 40% reduction. Coupled with the amount of performance the company manages to wring from the silicon, AMD says the lower power consumption leads to a 22% gain in power efficiency (perf-per-watt).

The £269 Ryzen 5 9600X delivers strong processing at its price point. At stock settings, the £300 Core i5-14600K is 4% faster in 1080p gaming. Bear in mind that the Ryzen 5 9600X delivers these results from within a 65W TDP envelope.

The Ryzen 5 9600X is 12% faster than the previous-gen Ryzen 5 7600X, a solid generational gain, especially considering the 7600X's 105W TDP. Enabling the auto-overclocking PBO feature and bumping memory speed up to DDR5-6000 yields a solid 8% increase in performance, but it's possible that more directed tuning could expose larger increases.

On the gaming front, AMD's stiffest competition comes from its own roster. The Ryzen 7 5800X3D is 11% faster than the 9600X at 1080p gaming for a £50 premium, while the Ryzen 7 5700X3D offers the same performance as the 9600X on average but for £65 less. It also comes on the cheaper AM4 platform. These chips do come with the older Zen 3 architecture, however, so they won't be as performant in daily tasks.

The Ryzen 5 9600X delivers solid generational gains in single-threaded performance, matching or exceeding the price-comparable Intel models. It is 8% faster in our cumulative measure of single-threaded performance than the prior-gen Ryzen 5 7600X and 4% faster than the Core i5-14600K.

The Intel processors manage to keep their lead in our cumulative measure of heavily threaded productivity application performance, and often by large margins, but there are a few caveats. The Core i5-14600K is 26% faster than the Ryzen 5 9600X. In



The most affordable Zen 5 processor is also the most attractive right now.

threaded work, the Ryzen 5 9600X is 14% faster than the prior-gen 7600X. Again, that's while consuming quite a bit less power, so the generational efficiency gains are even more impressive.

While the Ryzen 7 5800X3D and 5700X3D are exceptional in our gaming benchmarks, rivalling the Ryzen 9000 chips at a lower price point, pay attention to the huge disparities in application performance with these Zen 3 chips. The 9600X is also 18% faster in threaded workloads than the 5800X3D.

The Ryzen 5 9600X sets a new bar in the value segment. Its solid gaming and single-threaded performance, coupled with exceptional power efficiency and a reasonable price, earns it a spot on our list of the best CPUs for gaming. Intel's competing Core i7-14600K still holds an advantage in heavily threaded workloads, but AMD has a commanding lead in gaming that helps offset the advantage. Additionally, the Ryzen 5 9600X's stellar power consumption metrics will ultimately yield a quieter system that has more forgiving cooling and power requirements.

## VERDICT

**DEVELOPER:** AMD

**WEB:** [www.amd.com](http://www.amd.com)

**PRICE:** £269

<b>FEATURES</b>	<b>9/10</b>	<b>EASE OF USE</b>	<b>9/10</b>
<b>PERFORMANCE</b>	<b>8/10</b>	<b>VALUE</b>	<b>9/10</b>

Sets the new bar for mid-range gaming systems built around standard processors. Solid gaming performance and single-threaded performance, coupled with exceptional power efficiency.

» **Rating 9/10**



# AMD Ryzen 9 9950X

If you can afford this, good for you!

## SPECS

**Socket:** AM5  
**Arch:** Zen 5  
**Process:** 4nm TSMC CPU, 6nm I/O  
**Cores:** 16  
**Threads:** 32  
**Clock:** 4.3GHz (5.7GHz boost)  
**Cache:** 1280KB L1, 16MB L2, 64MB L3  
**Unlocked:** Yes  
**GPU:** AMD Radeon 2 Cores  
**Mem:** DDR5-5600, ECC mobo support, two-channel, 192GB max  
**PCIe:** v5 28-lanes  
**TDP:** 170W

**A** MD's Ryzen 9 9950X comes with 16 cores and 32 threads, slotting in as the flagship model for the company's new Granite Ridge family of processors sporting the Zen 5 architecture. AMD has infused its innovative thread-targeting technology, previously reserved for the gaming-optimised 3D V-Cache models, into the upper-tier Ryzen 9000 models. However, the improvements in gaming performance aren't enough to take the crown from Intel's competing Core i9-14900K flagship, let alone AMD's own X3D processors.

Unlike Intel's competing Raptor Lake processors, Ryzen 9000 also has full native AVX-512 support, a boon for productivity work. The 9950X also delivers a notable gain in single-threaded performance, significantly reducing the gap with the 14900K.

AMD's Ryzen 9 9950X takes the lead in our suite of multi-threaded application tests. While it delivers some uplift in gaming, it relies on the strength of performance in productivity workloads to justify its price tag, especially because stepping up to higher-end workstation-class fare like AMD's Threadripper Pro and Intel's Xeon W has become exceedingly expensive.

## Ryzen up to the challenge of our rival

The Ryzen 9 9950X replaces the previous-gen Ryzen 9 7950X, but aside from the newer CPU microarchitecture it has the same basic accommodations, such as 16 cores, 32 threads, a 5.7GHz peak frequency, 80MB of combined L2+L3 cache, and a 170W/230W TDP/max rating. AMD improved the 9950X's memory support from DDR5-5200 to DDR5-5600 and expanded the L1 data cache (L1D) to 48KB. AMD supports ECC memory with its silicon, but the implementation, validation and support are up to the motherboard OEMs.

The immediate takeaway is that if you're looking for the fastest gaming chips on the market, look no further than AMD's previous-generation Ryzen 7000X3D models. AMD has Ryzen 9000X3D processors slated for release in the future, so those in pursuit of the leading edge might be best served by waiting. In our measure of 1080p gaming, the Ryzen 9 9950X was approximately 8% faster than the prior-gen Ryzen 9 7950X. The competition is stiff, though. The £450 Core i9-14900K is 10% faster in 1080p than the £594 Ryzen 9 9950X, while the £380 Core i7-14700K is 5% faster.

At stock settings, the Ryzen 9 9950X is 12% faster than the previous-gen Ryzen 9 7950X in our multi-threaded performance testing, and that gain extends to 16% after we enable PBO and faster EXPO memory profiles for both processors.

The Ryzen 9 9950X is also 23% faster than the Core i9-14900K in heavily threaded work, an impressive result. We factor in performance with the



If throwing money is your aim, then there are better options at the moment, even from AMD.

AVX-heavy *y-cruncher* benchmark for this metric, so the 9950X's full support for native AVX-512 pays big dividends in our overall ranking. Without that benchmark factored in, the 9950X is still 13% faster than the 14900K, partly due to other applications that also leverage AVX instructions.

The Ryzen 9 9950X has the same 170/230W TDP rating as its prior-gen counterpart, the 7950X, but it does consume 7-15% less power in a range of workloads, such as *y-cruncher*, *HandBrake* and a few of the *Blender* renders. The 9950X delivers more performance than its predecessor with less power, but the efficiency gains will vary by workload.

The baffling situation is that most users focused on productivity work would be better served by a reduced-cost Ryzen 9 7950X – this chip will surely see further price reductions in the near future – or an Intel alternative, if the price is right. The Ryzen 9 9950X also doesn't make a compelling choice for gaming – AMD's own X3D processors remain the chips to beat without the extra platform costs. **LXF**

## VERDICT

**DEVELOPER:** AMD  
**WEB:** [www.amd.com](http://www.amd.com)  
**LICENCE:** PRICE: £594

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	7/10	VALUE	5/10

Offers performance improvements across the board and the highest performance available on a mainstream PC platform in multi-threaded workloads. However, it lags behind competing chips in gaming, and the generational gains are small enough in some productivity workloads that the previous-gen Ryzen 9 7950X is an attractive alternative.

» **Rating 6/10**

# Easy de-archiving

**Shashank Sharma** considers himself to be a multitasking maestro, and is naturally impressed with nifty utilities that can do the job of many.



**OUR  
EXPERT**

**Shashank Sharma** is a trial lawyer in Delhi and an avid Arch user. He's been writing about open source software for 20 years, and lawyering for over 10.

### QUICK TIP

You'll have to install the Python package `argcomplete` if you want to use auto-complete for *Patool's* subcommands. *Patool's* documentation on the subject is insufficient, and not applicable for all distros. A quick web search will provide answers.

**T**he Linux distro is home to a large number of compression utilities, each with its own myriad subcommands and options. Remembering them all is a Herculean challenge, even for veterans. What complicates the situation further is that no single tool is equipped to handle all the different file formats, such as TAR, BZ2, GZIP, GZ, TAR.GZ, RAR and so on.

Released under the GPLv3 licence, *Patool* is a portable archive file manager written in Python. You no longer have to rely on separate commands or remember their myriad options, as *Patool* supports just about all archive and compression formats. Even better, it uses sane and generic command options such as `extract` and `create` to respectively decompress an archive or to make one from the supplied set of files.

The project's website suggests running the `sudo pip install patool` command to install *Patool*. While this might work for some distros, Debian and recent iterations of Ubuntu no longer allow installing Python packages in this way. If your distro throws up an `externally managed environment` error when you run the `pip install` command, but you still want to install packages using *Pip*, your only option is to create a virtual environment. Thankfully, there's a quicker workaround that works for all distros.

You must have Python 3.10 or greater on your machine, as all other dependencies are likely already installed if you've been using your distro for any length of time. Head over to the project's GitHub page (<https://github.com/wummel/patool>) and download the latest tarball. You can now use the bundled `setup.py` script to install *Patool*:

```
$ tar xvf patool-3.0.1.tar.gz
$ cd patool-3.0.1/
$ sudo python setup.py install
```

The last line utilises the `setup.py` script to install *Patool* in the `/usr/local/bin/` directory.

### Archives demystified

The objective of *Patool* is to simplify the process of working with archive files. You can use the nifty utility to create archives, extract files, list archive contents, compare archives, search for files within archives, and even repack archives to a different format.

First, let's create an archive with `patool create test-archive.zip <filename1> <filename2>`:

```
1: linuxlala@playground: /media/linuxlala/Stuffsies/distros
1: linuxlala@playground: /media/linuxlala/Stuffsies/distros
$ patool create test-archive.zip BCD_Verification.pdf Letter-delay.pdf
INFO patool: Creating test-archive.zip ...
INFO patool: ... test-archive.zip created.
$ patool list test-archive.zip
INFO patool: Listing test-archive.zip ...
BCD_Verification.pdf
Letter-delay.pdf
```

No matter the archive format, *Patool* always extracts the contents into a dedicated directory.

```
$ patool create test-archive.zip BCD_Verification.pdf
Letter-delay.pdf
INFO patool: Creating test-archive.zip ...
INFO patool: ... test-archive.zip created.
$ patool list test-archive.zip
INFO patool: Listing test-archive.zip ...
BCD_Verification.pdf
Letter-delay.pdf
```

*Patool* automatically identifies the type of archive you wish to create, and runs the relevant command to perform the operation. For instance, when creating a TAR.GZ archive, *Patool* runs the `usr/bin/tar --create --gzip --force-local --file another-archive.tar.gz --<filename1> <filename2>` command. It similarly runs the `usr/bin/tar --create --bzip2 --force-local --file third-archive.tar.bz2 --<filename> <filename-2>` command if you opt to create a TAR.BZ2 archive instead.

The `patool list <archive-name>` command can list the contents of an archive, without extracting files.

Another significant advantage of *Patool* is that a single command can be used to extract files from different archives. For instance, if you have a TAR.GZ and a RAR archive, instead of running separate commands to extract the files, you can run the `patool extract <archive1> <archive2>` command, where the two or more archives are different file formats, and *Patool* extracts the files.

The command `patool extract archive.zip another-archive.tar.gz third-archive.tar.bz2` extracts the files



from each archive into the **archive**, **another-archive** and **third-archive** directories respectively.

When extracting files from an archive, *Patool* is sensitive enough to not populate the contents of the archive in the current working directory. Instead, it extracts files from an archive to its own directory. This is especially useful if you use *Patool* to simultaneously extract files from multiple archives.

In addition to standard archive files, you can also use the `patool list` command to view the contents of DEB package and even ISO and CHM files, if you have the right packages installed.

You can run the `patool formats` command for a list of all the file formats supported by the utility. While most commonly used formats are supported out of the box, the command informs you of the packages you need to install if you want to work with other archive formats. For instance, we had to install the **7zip** package before we could work with 7Z archives.

It's also possible to use *Patool* to repack archives in one format to another with the `patool repack <source-archive> <destination-archive>` command:

```
$ patool repack test-archive.zip test-archive.tar.gz
INFO patool: Repacking test-archive.zip to test-archive.tar.gz ...
INFO patool: ... test-archive.zip extracted to '/tmp/Unpack_wpq74mci'.
INFO patool: running /usr/bin/tar --create --gzip --force-local --file /home/linuxlala/Downloads/test-archive.tar.gz -- "Letter-delay.pdf" BCD_Verification.pdf
INFO patool: with input=""
INFO patool: ... repacking successful.
```

The `repack` command automatically identifies the source and target archive formats, and executes the relevant commands to create the new archive. The command doesn't replace the source archive, however, so with the above command, we end up with a **test-archive.zip** as well as **test-archive.tar.gz** files.

## Search and diff

*Patool* also supports searching for files within archives. The archive is first extracted into a temporary directory and then the `grep` utility is used to find matching files. Unfortunately, we found the feature to be unreliable during our tests.

The command `patool search "string" <archive-name>` is straightforward enough, however the results are inconsistent. While the project could easily identify files in ZIP archives, *Patool* struggled to return results when working with any other archive format:

```
$ patool search "index" third-archive.tar.bz2
INFO patool: Searching 'index' in third-archive.tar.bz2
...
INFO patool: running /usr/bin/grep -r -e index .
INFO patool: with cwd='/tmp/Unpack_e2rogteo', input=""
INFO patool: ... 'index' not found
$ patool list third-archive.tar.bz2
INFO patool: Listing third-archive.tar.bz2 ...
INFO patool: running /usr/bin/tar --list --bzip2 --force-local --file third-archive.tar.bz2
INFO patool: with input=""
```

```
linuxlala@playground:/media/linuxlala/Stuffsies/from-bricked-Ubuntu-20.04/spicejet-temp$ patool repack archive.zip archive.tar.gz
INFO patool: Repacking archive.zip to archive.tar.gz ...
INFO patool: running /usr/bin/7z x -o/tmp/Unpack_pxf59efu -- archive.zip
INFO patool: with input=""
INFO patool: ... archive.zip extracted to '/tmp/Unpack_pxf59efu'.
INFO patool: running /usr/bin/tar --create --gzip --force-local --file /media/linuxlala/Stuffsies/from-bricked-Ubuntu-20.04/spicejet-temp/archive.tar.gz -- 8-letter-dated-02-05-2019-rejecting-RTI-appeal.pdf 4-certificates-of-training.pdf 5-reply-to-RTI-dated-28-09-2018.pdf 3-baggage-tags.pdf 6-termination-order.pdf "JS M-vs-UoI-&-Ors-e-inspection-court-file -- password is jmy.pdf" 1-letter-of-offer.pdf 7-letter-dated-14.12.2018-rejecting-RTI-request.pdf 2-letter-dated-21-03-2016.pdf
INFO patool: with input=""
INFO patool: ... repacking successful.
linuxlala@playground:/media/linuxlala/Stuffsies/from-bricked-Ubuntu-20.04/spicejet-temp$ patool diff archive.zip archive.tar.gz
INFO patool: Comparing archive.zip with archive.tar.gz ...
INFO patool: running /usr/bin/diff -urN /tmp/Unpack_t8xh37t2 /tmp/Unpack_zeyj10g
INFO patool: with input=""
INFO patool: ... no differences found.
```

indexed-vakalatnama.pdf

index-vakaltanma.doc

As you can see, we in fact have two files in the archive that match the provided search string, but *Patool* failed to find them. Providing an exact filename didn't help either.

The `patool diff <archive1> <archive2>` command can be used to list differences in the content of specified archives. The archives can be in different formats, and you can even view the differences between more than two archives. The contents of the archives are extracted into a temporary directory but *Patool* doesn't inform you which temporary directory is home to which archive. The output is thus difficult to make sense of.

Despite its poor handling of the search and diff features, *Patool's* ease of use makes it an easy recommendation. However, the project could do with a little more polish and improved documentation. **LXF**

Make sure to delete the source archives if you use the 'patool repack' command to create a different archive format.

## » THE PERL, PYTHON AND RUST PALS

While we often run into projects that provide similar functionality, most of them are usually programmed in the same language. This is what makes *Atool*, *Patool* and *Ouch* (<https://github.com/ouch-org/ouch>) unique.

Written in Perl, *Atool* ([www.nongnu.org/atool](http://www.nongnu.org/atool)) is the oldest of the three, first released in 2001. It supports many popular archive formats, but wasn't featured as the subject of this tutorial because it hasn't seen a new release since 2012.

Still, *Atool* boasts of all of *Patool's* features, such as extracting files from an archive into a dedicated directory, repacking archives into a different format, and so on. One key difference between the two tools, however, is that while *Patool* uses subcommands for various operations, *Atool* comprises several different tools for each task it can perform. These include `aunpack` to extract files from an archive, `apack` to create archives and `als` to list the contents of an archive. You'll find *Atool* in the software repositories of most desktop distros, so installing it is only a matter of running the `sudo apt install atool` command on DEB-based distros or `sudo dnf install atool` command on RPM distros.

*Ouch* (*Obvious Unified Compression Helper*) is a Rust utility that offers the same functionality as the Perl and Python counterparts. You use the `ouch decompress` command to extract the contents of an archive. The `ouch compress` command can similarly be used to create an archive. Run the `cargo install ouch` command to install it.

» **ENHANCE YOUR TERMINAL-FU** Subscribe now at <http://bit.ly/LinuxFormat>

## LINUX BASICS

**Part Eight!**  
Don't miss  
next issue,  
subscribe on  
page 16!

# Taking control of hardware drivers

**Nick Peers** discovers how to use your distro and third-party tools to help profile your hardware, find missing drivers and more...



**OUR EXPERT**

**Nick Peers** is grateful for the efforts of thousands of developers and enthusiasts in ensuring most hardware works without a hitch on his PC.

**O**ne of the most critical components in your Linux installation is its collection of drivers. These are responsible for recognising all the hardware connected to your PC, from internal components like motherboards, network adaptors and graphics cards to external peripherals such as your all-in-one device and USB Wi-Fi adaptor. Without the correct drivers in place, your PC can't communicate with these components, rendering them useless.

Once upon a time, getting drivers working on your Linux PC was a real pain, but these days, thanks to the efforts of developers, enthusiasts and – yes, we'll admit it – hardware manufacturers, it's possible to install a fresh version of Linux on your PC and never have to worry about plugging anything in for the first time. Yep, plug and play is a real thing in today's Linux distros.

But like all OSes, things aren't always perfect, with obscure chipsets and obsolete hardware capable of causing problems. And that's where we come in...

Drivers are part of the kernel, the core Linux system, and bridge your hardware to your OS. Like everything

else in the Linux filesystem, hardware devices are treated as files, and these can be viewed in the `/dev` directory. There are three broad types of device driver: character, block and network. They refer to how data is read and written – examples of each type include keyboard and mice (character), any storage device, like a hard drive (block), and your Wi-Fi adaptor (network).

As with Windows and Mac OS, many hardware drivers are bundled with the OS itself – in the case of Linux, they're included as part of the kernel. Each kernel update not only updates existing drivers, but adds support for the latest hardware devices, too.

## Identify hardware

In an ideal world, your PC should know everything about its hardware components, including any devices attached to it. But how do you confirm this and find out what hardware has been recognised by your system?

One option is to trawl through your distro's *Settings* app – we covered both Ubuntu and Mint in **LXF318's** *Basics* column. But if you'd like a handy summary of


## QUICK TIP

Open the **Launcher** and search for 'firmware' to launch the **Firmware Updater** tool. This checks key hardware components – including hard drives – that contain upgradeable firmware to see if any updates are available. If they are, you can use the tool to update them safely.

## » OPTIONS FOR GRAPHICS DRIVER

Unlike many devices, graphics chips are constantly in flux as new features, optimisations and bug fixes are developed. As a result, new graphics drivers are released regularly by the manufacturer, and Linux is well supported by all three major manufacturers: AMD, NVIDIA and Intel.

Compatible drivers for all three as well as a host of other chipsets are built into the kernel (see [www.kernel.org/doc/html/latest/gpu/drivers.html](http://www.kernel.org/doc/html/latest/gpu/drivers.html)). But there are other options if you wish – particularly if you're looking for more features or better gameplay – by installing alternative drivers to the default on offer.



Nvidia's open kernel modules receive the most testing, and have almost reached feature parity with non-open modules.

Both Ubuntu and Mint make this choice easy – open Software & Updates from the Launcher (or Driver Manager in Mint) and switch to the Additional Drivers tab. Here you'll see whether any alternative drivers are recommended – those with

Nvidia graphics chips are offered a range of Nvidia drivers (both open kernel and proprietary) alongside the default Nouveau display driver, for example. Pick a different version and click Apply Changes to upgrade or switch.

With AMD cards, you can stick with the supplied open source kernel drivers, investigate the open source AMD Vulkan drivers (<https://github.com/GPUOpen-Drivers/AMDVLK>) or try your luck with AMD's proprietary drivers ([www.amd.com/en/support/download/linux-drivers.html](http://www.amd.com/en/support/download/linux-drivers.html)) in the hope they've been updated to support Ubuntu 24.04 – at the time of writing, they hadn't been.



key hardware, one tool places it all at your fingertips: *System Profiler & Benchmark*. Install it via the terminal:

```
$ sudo apt update && sudo apt install hardinfo
```

The annotation (*right*) reveals how to use this to navigate your hardware and glean all kinds of useful data about how it's performing. One particularly handy section is *Boots*, which displays a list of boot times and which kernel version was used. As we'll see, this can be helpful when determining if a new driver installed as part of a kernel update is the cause of your problem.

## Find hardware drivers

In an ideal world, your hardware just works. On rare occasions, you plug something in and nothing happens. This often indicates that drivers aren't available for that device in the kernel, which means you need to source and install them separately.

First check any documentation that came with the device – if it doesn't provide explicit instructions for installing drivers in Linux, it may give you details about the adaptor's make and model, and a web address to visit. Here, look in the Support section for downloads or drivers, and see if there's information for Linux.

In most cases, drivers are supplied in tar or Gzip format; the instructions should provide the commands you need (usually issued via a terminal window) to unzip the archive and install the driver, typically through running a *Bash* script (a file with a .sh extension).

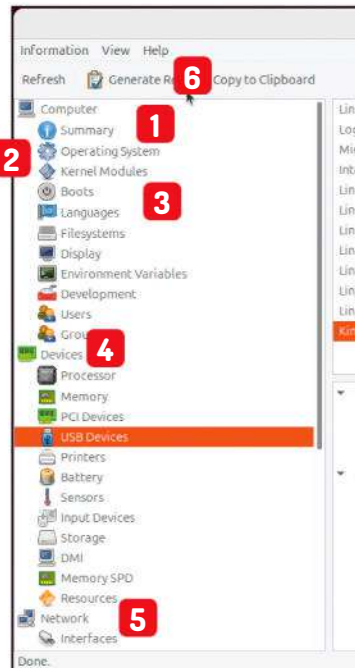
If you don't have access to any identifying information about the hardware, there's a good chance you can get it via *System Profiler*. Open the app and explore the various sections to see if the hardware has been detected somewhere in the system.

For example, one common issue people encounter is with Wi-Fi adaptors, so here you'd first look under *Network > Interfaces*, where it should hopefully be listed alongside your other network adaptors, such as your wired Ethernet adaptor. This indicates the hardware has been detected and correctly identified at least, but doesn't reveal much information about it.

To get more info, go to *Devices > PCI Devices* if it's an internal plug-in card, or *Devices > USB Devices* if it's an external USB adaptor to see if you can identify it (look for a Network Controller entry, which should reveal the name and chipset of the Wi-Fi adaptor).

Armed with this information, you can then source potential drivers online. Start with official sources – namely, the manufacturer's website. If this fails, widen

## PROFILE YOUR SYSTEM



- 1 Summary**  
Displays a simple list of all recognised hardware, from internal components to external peripherals.
- 2 Kernel Modules**  
Provides a comprehensive list of device drivers, including name, description, version number and author.
- 3 Boots**  
Displays a handy summary of which kernel version was booted when – can help pinpoint problems to recent kernel updates.
- 4 Devices**  
Provides details about specific hardware types – the PCI and USB device sections are particularly informative.
- 5 Network**  
A series of networking tools – Interfaces has details of wired and wireless connections.
- 6 Generate Report**  
Create reports in plain text or HTML format from the Computer, Devices, Network and Benchmarks sections.

your web search to a more generic 'linux' plus the device or chipset name. If too many conflicting results are returned, try replacing 'linux' with your distro (such as Ubuntu or Mint) to see what comes up. If you're searching for better graphics card drivers, check the box (*opposite page*) for details on what's available.

When it comes to other drivers, in some cases you may strike lucky and reveal links to official or unofficial drivers, or you may come across a slew of forum posts from other people trying to source Linux drivers. Read through the entire threads – the solution is often found in there somewhere, or they might reveal further bits of information you can use to modify your web search.

## Troubleshoot problems

Most hardware problems can be traced to drivers, and those drivers have usually recently been installed or updated following a kernel update. As with most computing problems, taking precautions makes life so much easier. The step-by-step guide (*page 61*) reveals how to use *Timeshift* to start taking system snapshots, enabling you to easily roll back recent driver installations or kernel updates to restore a working system again.

However, occasionally hardware can stop working for no apparent reason. How can you determine what's wrong? Start by checking if its driver is loaded. First, you need to identify what driver (or more accurately, what kernel module) it's using. This requires access to the terminal. First, identify the hardware itself:

```
$ sudo lspci -v
```

Locate your malfunctioning hardware in the list and look for the lines marked 'kernel driver in use' (indicating it's loaded and running) and 'kernel module'. If you're troubleshooting a device attached via USB, use `sudo lsusb -v` instead.

If the kernel module is listed but not marked as in use, try loading the driver manually with *modprobe* to see if it gets things working again (substitute

### QUICK TIP

If you enjoy a simple life only attempt to use the graphic drivers offered by your distro. These are compiled and signed so will install with little issue. Trying graphic drivers from other sources is a recipe for blank screens.



Your Settings app should help you work out what hardware is in your system, but *System Profiler & Benchmark* has everything in one place.



## QUICK TIP

If you can't source a driver for your network printer, don't forget you can usually administer it through your web browser, which enables you to keep an eye on toner and ink levels. Check your printer manual for details on accessing it – typically through its IP address.

drivername with the full name of the driver as listed by *lspci* or *lsusb*):

```
$ sudo modprobe drivername
```

Confirm it's loaded with:

```
$ lsmod | grep drivername
```

## Troubleshoot using system logs

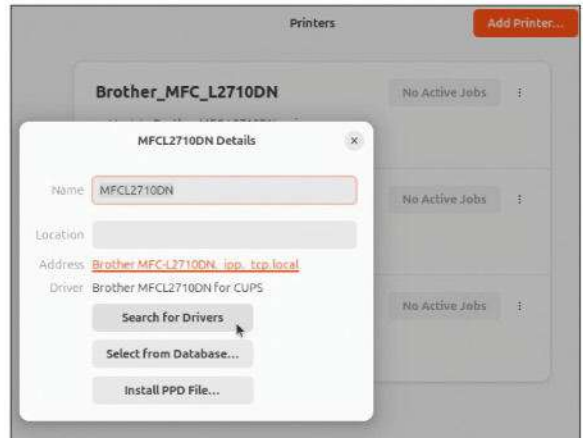
Another way to troubleshoot hardware problems – including those that don't appear to be the result of missing or unloaded drivers – is with system logs. The easiest way to view these is through Logs, a desktop tool you can access via Launcher (it's in Utilities).

If you're lucky, you'll see some signs of your problem in the Important view; if not, select Hardware on the left and either scroll through the logs looking for some telltale errors or messages, or use the search button to narrow your search using keywords such as the device or driver name (note, searches are case-sensitive).

When you spot a potentially useful message, click to highlight it. A pop-up appears displaying the message again; select this, right-click it and choose Copy, then paste the text into a web search window to see if you can find help online. Focus on driver names and specific error codes or messages if a search for the entire error message doesn't help. And if you're still unable to make head nor tail of a problem, consider emailing Neil at [answers@linuxformat.com](mailto:answers@linuxformat.com) with details to see if he can track down a possible solution.

## Printers made easy

One device that has become ever easier to set up and use is a printer. These days, most printers work simply by connecting them to your PC; if you have a network printer, they're often already up and running in the background. This is all thanks to the *Common Unix Printing System*, or *CUPS* for short. *CUPS* is Linux's tool for managing print jobs, queues and network



You can unlock all of your printer's advanced features by installing the relevant PPD file from the manufacturer's website.

printing. It's also capable of detecting and using any printer that has some form of Internet Printing Protocol (IPP) support, such as AirPrint or Wi-Fi Direct Print; 98% of printers currently sold support one of these protocols, which is how they work out of the box.

One drawback is that IPP support often means some advanced functions of your printer aren't supported. To gain access to your printer's full set of features, you may have to install a PostScript Printer Definition (PPD) file. Visit your printer manufacturer's website and look in its Support or Drivers section for such a file – if you have a multi-function printer with scanning capabilities, you may be offered a utility that downloads and installs all the required files to support printing and scanning, but if you're only offered a PPD file, save it to your **Downloads** folder.

If downloading a PPD file, open *Settings* in Ubuntu and visit the Printers section. Click the vertical ellipsis next to No Active Jobs and choose Printer Details, then click Install PPD File to install your driver.

Once installed, you may find there are two instances of your printer in the Printers section – this is normal, and you should be able to differentiate your newly installed driver from the driverless *CUPS* version by opening its Properties or Options, where you should find more options like toner save or image quality.

## Printer Applications

To make setting up and using printers easier, *CUPS* is evolving to make PPDs redundant through the use of Printer Applications. These aim to replace the need to source individual PPD files with apps that come with support for hundreds or thousands of models. One such printer app already deployed in Ubuntu is the *HP Linux Imaging and Printing (HPLIP)* app. This supports over 3,000 HP printers, scanners and fax machines.

Other printer apps offer different kinds of support – for example, *Gutenprint* (<https://gimp-print.sourceforge.io>) is designed to provide better quality output for various brands of photo printers. Once installed ( `sudo apt install printer-driver-gutenprint` ), choose Printer Details from Settings > Printers as before, but this time choose Select From Database. Look for an entry referencing both your printer model and 'Gutenprint'. Select this to switch to that driver.

One complication for printing is the presence of Snap in Ubuntu. Snap-based applications are

## » CUPS RUNNETH OVER

There are two versions of *CUPS* installed in the latest versions of Ubuntu. There's a native version, which is what you've been using so far, plus a *CUPS Snap* version, which is used as a proxy to relay print jobs from sandboxed Snap-based apps to the native *CUPS*.

The native version is built on the legacy 2.4.x branch of *CUPS*, which is the last version to support PPD files for printer drivers. *CUPS Snap* is built in *CUPS 3.0*, which aims to phase out PPDs in favour of Printer Applications. These act as intermediaries between driverless IPP printing and non-IPP printers – for example, *ipp-usb* exists to extend IPP capabilities to your network printer when you connect it to directly to your PC via USB cable. Another key Printer Application is *PS-printer-app*, which works with thousands of PostScript printers – many of them older models – out of the box.

One final Printer Application worth mentioning is the *Legacy Printer App*. By default, *CUPS Snap* doesn't work with PPD files, so even after installation, Snap-based apps won't recognise them. Thankfully, there's a workaround that enables you to tap into your printer's advanced features: *Legacy Printer App* makes your printer's PPDs visible to Snap apps. Unlike other printer apps, you can't install it through Snap; instead, open a terminal and issue the following:

```
$ sudo apt install legacy-printer-app
```

Once installed, you should find your Snap applications have access to the same printer functions as your non-Snap programs.



sandboxed, so need extra help communicating with your printer – see the box (opposite page) for details.


### Set up your scanner

If you have a scanner or all-in-one device, before attempting to install anything, open *Document Scanner* in Ubuntu and wait while it searches for scanners. If yours shows up, you're good to go. If it doesn't, first visit [www.sane-project.org](http://www.sane-project.org) and click Supported Devices to browse a database of known models and whether or not they're supported. If they

aren't, you may be given some advice for further research; if yours is supported, check your manufacturer's website for the required drivers that should allow it to be detected going forward.

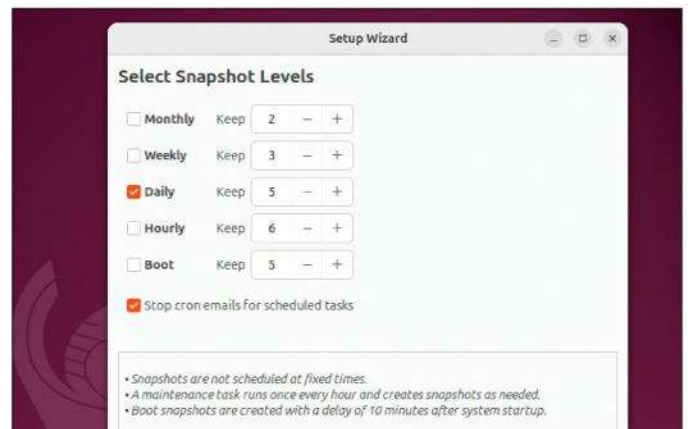
If you're unable to find any drivers for your scanner – particularly older models – then all may not be lost. Visit [www.hamrick.com/alternate-versions.html](http://www.hamrick.com/alternate-versions.html) to download a free trial of VueScan, which has reverse-engineered drivers for over 7,400 models so they'll continue to work long after the manufacturer has abandoned them. **LXF**

## CREATE A DRIVER ROLLBACK SYSTEM



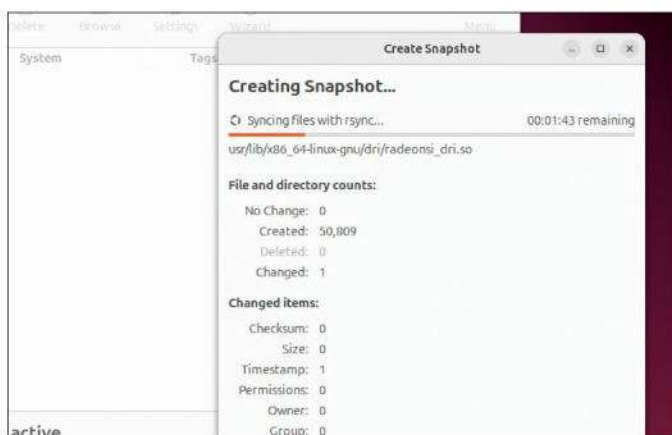
### 1 Set up Timeshift

*Timeshift* (or *System Snapshots*) is pre-installed with Linux Mint; Ubuntu users can install it via the terminal with `sudo apt install timeshift`. Once done, open it from the Launcher or Start menu. If given a choice of snapshot type, leave *Rsync* installed and click Next. *Timeshift* tries to locate a backup location – only Linux filesystems are supported. Select it and click Next again.




### 2 Set snapshot levels

The default setting checks hourly for changes, taking snapshots as and when needed. It keeps these for 24 hours, plus five daily snapshots, which should be sufficient for driver rollbacks, but you can alter these settings if you wish. Click Next when you're done, leave all home directories excluded, and click Next followed by Finish to be taken to the main screen.



### 3 Taking snapshots

Click Create to take your first snapshot. Once taken, it appears in the main *Timeshift* window, while future snapshots are taken automatically behind the scenes. If you're about to upgrade your kernel and/or a driver, you can re-open *Timeshift* to take a manual snapshot of your system before installing it, instead of relying on the most recent automatic backup.



### 4 Rolling back drivers

If you encounter a hardware problem after installing the update or driver, open *Timeshift*, select the snapshot taken before it was applied and click Restore. Follow the prompts, leaving the default settings in place. Review what files will be restored or removed, then click Next. After restoration, reboot to complete the rollback and remove the errant file.

» **IMPROVE YOUR LINUX SKILLS** Subscribe now at <http://bit.ly/LinuxFormat>

## TAILSCALE

Credit: <https://tailscale.com>

# Use Tailscale to create a mesh VPN

Craving constant connections, **Matt Holder** discovers how to use Tailscale to create a mesh network between all of his devices.



**OUR  
EXPERT**

**Matt Holder** is an IT professional of 15 years, Linux user for over 20 years, user of plenty of home-automation gear and self-professed geek.

**W**e are going to investigate how we can use VPNs to keep our home networks safe, while being able to access resources when we are away from home. We are also considering how to use a VPN to provide remote support while not opening up all the devices to which we have access to all users.

Before getting into specifics, let's talk about what a VPN (virtual private network) is. A VPN allows us to expand our networks, while providing a virtual network layer. It is often used to bridge two networks at remote locations together or allow access to resources in a more secure manner than opening ports on a firewall.

There is a large number options available (OpenVPN, SSL VPN, ZeroTier and many others) and these operate at Levels 2, 3 or 4 of the OSI model, with some creating a star topology back to a central location and some operating in a peer-to-peer manner. In this article, we are focusing on the commercial Tailscale VPN, which has a generous free tier available.

Tailscale operates by providing a peer-to-peer network between your devices, where none of your traffic goes through its servers, but they are used to provide the control plane, including negotiation services, to initiate the connections. Tailscale refers to your private network as a Tailnet. Management of your Tailnet and Tailscale account is through its web portal and from here you can configure everything, including reauthentication times for devices, DNS settings and access control lists (ACLs).

Tailscale uses the modern and well-regarded (and secure) WireGuard VPN technology, and provides methods of key exchange and all of the other features that will be discussed.

So, with that little bit of introduction given, what will we do in this article? First of all, we will detail the creation of an account and download the client for one of your devices. While talking about devices, clients are available for Windows, Mac OS, iOS, Android and Linux. Once the account has been created and the first device added, we will investigate the admin portal a little further to demonstrate some of the key principles. Following this, we will spend some time talking about how you could use the Tailscale add-on for *Home Assistant*, which will allow access to your smart home kit without opening ports on your router. We will then

spend some time introducing some of Tailscale's additional functionality.

## Usage scenarios

There are countless different things that Tailscale can be used for, with the first being access to your home network from external locations. This is more secure as it means that you do not need to open ports from the big wild internet to the safe haven of your home.

By adding the client to all your devices, you have an IP network wherever you are. This means you can do things like send files from worldwide locations to anyone in your household. *LocalSend* is perfect for this task and can run on most operating systems.

More advanced uses include allowing access to staging/production servers from a certain set of devices or users, while securing other devices. Similarly, only giving access to the finance package to finance staff would also be a good idea.

## Create your Tailnet

Before doing anything else, we need to create our account. Visit [www.tailscale.com](https://www.tailscale.com) and click on the Get Started button. On the page that loads, choose to sign in with a Google, Microsoft, GitHub or Apple account, or enter an email address. Once this has been done, follow the wizard to create your account and install the client on your first device.

To use even the most basic of features, you need to install the client on two or more devices. If you install the client on your phone and a device on your home network, you can connect between the devices wherever you are. If you are a *Home Assistant* user, head to the Add-On store to add the Tailscale VPN container. Once added and connected, you can connect to *Home Assistant* from outside of the network, without opening ports on your firewall.

In the rest of the article, we will discuss some of the more advanced features that can be used to get the most out of your new VPN.

## Mask your location

As Tailscale, by default, only traverses the traffic that is strictly necessary, this means that if you are using a public Wi-Fi network somewhere, you don't have the

### QUICK TIP

You can learn a lot more about Tailscale by visiting the excellent website here: <https://tailscale.com/kb/1017/install>



peace of mind that your internet traffic isn't being snooped. This is where Exit Nodes come in. This is where you can connect to Tailscale, enable the feature, and your internet traffic is carried over the VPN link. You could create Exit Nodes in different countries, which would enable you to appear to be browsing the web from different places across the world.

To set this up, navigate to the Machines section of your account and find the device you wish to configure as an Exit Node. Select the device and under the Route option, click on Edit and turn on the Exit Node option. Save the changes and next time you connect from your phone or laptop, you can turn on the Exit Node to enable you to send your internet traffic via the VPN and this device.

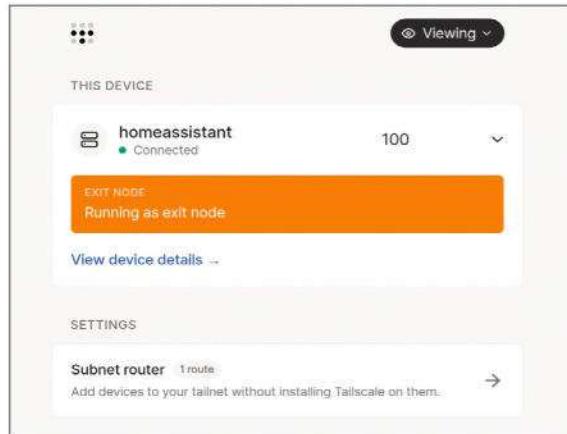
## Device support

By installing the VPN client on as many of your devices as possible, you can use Tailscale as effectively as possible. However, there are devices that can't run the client, such as printers, network switches and routers. This is where subnet routers come in. Pick a device in your home that is running the client and you can configure it to advertise your local subnet. This is a two-step process, first of all starting on your Linux box. Open a terminal and run the following:

```
$ echo 'net.ipv4.ip_forward = 1' | sudo tee -a /etc/sysctl.conf
$ echo 'net.ipv6.conf.all.forwarding = 1' | sudo tee -a /etc/sysctl.conf
$ sudo sysctl -p /etc/sysctl.conf
$ sudo tailscale up --advertise-routes=192.168.1.0/24,
```

In the listing above, the first two commands instruct the Linux kernel to allow IP forwarding, which means that traffic can flow between the VPN network interface and the interface connected to your network. The third command connects to your Tailnet and advertises your network IP addresses. Make sure you change this to your range. In the third command, the /24 represents the number of IP addresses that should be advertised. This range covers 192.168.1.1-192.168.1.254, so be sure to update this according to how your network is set up.

Finally, from the Machines section of your dashboard, select the device you've just configured, open the Route options and select to allow the range. Once this has been completed, when you connect to



The Home Assistant add-on enables you to connect to your smart home from outside of the house, without opening ports on your firewall. Note that this client is advertising a route that allows access to other devices that do not/cannot run the client.

the Tailnet from other devices, you will be able to connect to other IPs in your network.

## Send files

Taildrop is a brilliant feature, which can be used to send files between devices. It needs to be switched on from the General section of your Tailnet dashboard, then file sending options are available from the client on any of your devices.

On an Android device, find the file you wish to send, select the Share option and scroll down to find Tailscale from the intent list. Once selected, you see a list of your devices, which you can pick from and then the file is sent to the selected device.

In Windows, once the client has been installed, you can select a file in Windows Explorer and then when it has been right-clicked, a context menu item called Send With Tailscale is available. Using this enables you to select a device.

Using a Linux device requires you to send files from a command line. Open a terminal and use the following:

```
$ tailscale file cp /PATH/TO/FILE(S) DEVICE_NAME_OR_IP:
```

The syntax of the command has been designed to match that of the venerable *scp*.

## Serve and Funnel

These two features allow web-based services to be made available to other users. Serve allows users on your Tailnet to access advertised services as if they were available from the outside world, while Funnel is

## » HEADS-UP ON HEADSCALE

None of the traffic that flows between your devices flows through any of the Tailscale servers. If you still prefer not to trust a company with any of your VPN service, Headscale is the option for you. This is a project that takes the Tailscale open source code and bundles it together into a service that you can run yourself and use to connect your own clients to. To install Headscale on Ubuntu, open a terminal and run the following commands:

```
$ wget https://github.com/juanfont/headscale/releases/download/v0.23.0/headscale_0.23.0_linux_amd64.deb
$ sudo apt install ./headscale.deb
$ sudo systemctl enable headscale
$ nano /etc/headscale/config.yaml # Edit file accordingly
```

```
$ sudo systemctl start headscale
```

```
$ systemctl status headscale # Check headscale is running
```

Once Headscale has been installed, configured and is running, you can create your first user with the following command (running on the Headscale server):

```
$ headscale users create myfirstuser
```

The client can then be registered against your server by using the following command (running on the client):

```
$ tailscale up --login-server <YOUR_HEADSCALE_URL>
```

Finally, the machine needs to be registered on the Headscale server (running on the Headscale server):

```
$ headscale nodes register --user myfirstuser --key <YOUR_MACHINE_KEY>
```

**QUICK TIP**

An alternative to Tailscale is ZeroTier: <https://github.com/zerotier>

probably better considered as a temporary measure, where you could spin up a service and share it with a colleague who is not connected to the same Tailnet.

One of the excellent benefits of Tailscale Serve is being able to use Let's Encrypt to generate valid HTTPS certificates, to further secure users of your private services.

These features are quite advanced and due to the space available will not be covered further. See these links for more: <https://tailscale.com/kb/1223/funnel> and <https://tailscale.com/kb/1312/serve>.

**Tags and ACLs**

Tags can be used to differentiate between devices, their function, or the network they connect to. For example, a container running the open source audiobook management software, *Audiobookshelf* could be tagged as 'media', 'service' or both. An example of using tags to define a device's network would be to tag all devices, containers or services on a testing/staging environment as 'staging'. Devices and services that make up a production environment would be tagged as such. Devices can then be tagged to denote their purpose – for example, the devops team devices can be tagged as such and office staff can be tagged separately.

The tagging of resources is only the beginning of the story and they can be used in the hugely powerful ACLs to limit connections between resources. We will look at some examples below:

```
{
  "acls": [
    {
      "action": "accept",
      "src": ["tag:prod"],
      "dst": ["tag:prod:*"],
    },
  ],
}
```

The example above allows all resources tagged as production to connect to all other devices tagged as production. You could also segregate devices, so that web servers could connect to database servers, but other devices could not.

```
{
  "acls": [
    {
      "action": "accept",
      "src": ["group:devops"],
      "dst": ["tag:server:22"],
    },
  ],
}
```

In this example, a group of devices with the name devops are allowed to connect to any devices tagged as 'server'. The only port that can be connected to by the devops team is SSH. Here, the difference between **tag:devops** and **group:devops** is that the group contains a list of users, whereas the tag would contain a number of devices.

Combining what has already been discussed, we can see the next example:

```
{
  "acls": [
    {
```

```
      "action": "accept",
      "src": ["tag:webserver", "group:devops"],
      "dst": ["tag:database:*"],
    },
  ],
}
```

In this example, we can see that only devices tagged as 'webserver' and users in the devops group can connect to all ports on the database servers.

**Tailscale SSH**

SSH is a remote access solution for Linux machines. It replaced the venerable RSH, which was unencrypted. Typically, to use SSH you need a server on the machine you want to access and a client on the machine you are connecting from. Authentication can be carried out using passwords, encryption keys and certificates.

The tricky part of SSH, especially when operating at any sort of scale, is credential management. If you had to grant access to a bunch of new machines, you would need to generate the certificate files, copy the relevant key to each device and pass the other key to the user. This is made more complicated when a user no longer needs access – this now needs to be done in reverse, sometimes at great pace.

Tailscale SSH has been designed to solve this problem. By default users have access to Tailscale SSH and the service can be advertised on devices by running the following command in a terminal:

```
$ tailscale up --ssh
```

A policy needs to be added in the ACLs section:

```
{
  "action": "check", // "accept" or "check"
  "src": [list-of-sources],
  "dst": [list-of-destinations],
  "users": [list-of-ssh-users],
  "checkPeriod": "20h", // optional, only for check actions. default 12h
},
```

Using these policies, devices can be allowed access to others as well as controlling the users who can access devices. Should a user no longer need access, the rule can be removed. SSH recording can even be enabled if required. Using a check policy is especially powerful as it requires the user to re-authenticate after the period defined and is useful for securing root access. The **users** stanza defines which users the Tailscale user can connect as. For example, the **src** section can reference a user by their email address that is used to log in to Tailscale and the rule would allow access to all users in the **user** section.

Powering the SSH functionality is the client listening on port 22, acting as the server and performing the certificate copying with the Tailscale web service.

**Docker containment**

Tailscale provides its own containers, which allow services to be added directly to your Tailnet. This is carried out by using what is known as a sidecar container. Behind the scenes *Docker* is using the kernel namespacing functionality to merge the networks of the Tailscale container and whichever service you are running. This means your service is connected directly to your Tailnet and can be accessed by any devices that are registered. Don't forget, you can also use the



## » ENTERPRISE FEATURES

We have covered a large number of features, but there are many more on offer. For example, ACLs can be managed using GitOps, to allow a level of version control. GitHub actions exist, so that when a PR lands, the file is sent to Tailscale for validation. When files are merged into the main branch, the ACL file is synced up to Tailscale. Relatively recently, the Tailscale team has added support for device posture management. This means devices are only allowed to connect to resources when certain conditions are met. See below:

```
"postures": {
  "posture:latestMac": {
    "node:os == 'macos'",
    "node:osVersion == '13.4.0'",
```

```
"node:tsReleaseTrack == 'stable'",
  ]
},
"acls": [
  {
    // Only devs can access production and production access
    requires macOS is also up to date
    "action": "accept",
    "src": ["group:dev"],
    "dst": ["tag:production"],
    "srcPosture": ["posture:latestMac"]
  }
],
```

Access Control List functionality to limit which devices can connect to which services. An example of this functionality can be seen below:

```
---
version: "3.7"
services:
  tailscale-nginx:
    image: tailscale/tailscale:latest
    hostname: tailscale-nginx
    environment:
      - TS_AUTHKEY=tskey-client-notAReal-
        OAuthClientSecret1Atawk
      - TS_EXTRA_ARGS=--advertise-tags-tag:container
      - TS_STATE_DIR=/var/lib/tailscale
      - TS_USERSPACE=false
    volumes:
      - ${PWD}/tailscale-nginx/state:/var/lib/tailscale
      - /dev/net/tun:/dev/net/tun
    cap_add:
      - net_admin
      - sys_module
    restart: unless-stopped
  nginx:
    image: nginx
    depends_on:
      - tailscale-nginx
    network_mode: service:tailscale-nginx
```

We define two services that will be configured, with the first being the Tailscale container. The **environment** section contains authentication details, which are unique to your account, as well as a state directory, for files to be stored. The second service is nginx and the relevant lines here are **depends\_on**, which states nginx is dependent on Tailscale and the **network\_mode** line, which links the networking to the first container.

Create an **nginx** folder in your **home** directory and save the listing above as **compose.yaml**. Next, create a directory called **tailscale-nginx**, which will be mapped from the Docker container. Before starting the containers, you need to visit Settings > OAuth Clients in the Tailscale dashboard to generate a new oauth key. Copy the details that have been generated and replace the placeholder in the listing above, under the **TS\_AUTHKEY** section, before saving the file again.

Docker Compose can now be used to start the containers. Open a terminal, navigate to the **nginx** directory you created and run **docker compose up**.

If all is well the containers start and when you refresh the Machines listing in the portal, you see the new device. From a browser on a Tailnet connected device, browse to the IP address of the new container, as shown on the portal, and you should see a test page.

### More than a number

The MagicDNS feature enables devices to be referred to by a name rather than IP address. This could be **webserver.tail123abc.ts.net**. The suffix is the DNS name of your Tailnet and the prefix is taken from the name of the resource when it was first registered.

SplitDNS allows users to refer to devices using an easier-to-remember DNS name than the one provided. Access to the DNS provided by a local server is encrypted via the VPN tunnel, so can only be accessed by registered users. This also works when internal to the network or external and connected via VPN.

### Sharing

From the Machines page of your dashboard, you can share a device with the Tailnet of other users. This is really useful and allows you to provide remote support for a friend or family member. Only the shared device is available to the user you have sent the invite to. If the same user logs into multiple clients, they all have access. The shared device is quarantined so it can't initiate connections to devices in the other Tailnet.

To set up an ACL to stop the shared user accessing all services of the shared machine, you could use:

```
"acls": [
  // Admins can access everything on the tailnet.
  { "action": "accept", "src": ["group:admins"], "dst":
    ["*:*"] },
  // Shared users can only access port 80 and 443 of
  machines they are invited to.
  { "action": "accept", "src": ["autogroup:shared"], "dst":
    ["*:80,443"] },
]
```

We hope this quick tour of Tailscale has given you ideas about using your own mesh-based VPN. **LXF**

### QUICK TIP

If you wish to run your own control server, this can be accomplished by using Headscale: <https://github.com/juanfont/headscale>

» STOP CHASING YOUR OWN TAIL... Subscribe now at <http://bit.ly/LinuxFormat>

# BACK ISSUES » MISSED ONE?

## ISSUE 321 November 2024

Product code:  
**LXFDB0321**



### In the magazine

Build your own Linux fortress and protect your privacy with Qubes OS, and learn all about Matter, the new Internet of Things standard. Plus, discover which media server is best, find out how to secure your files, tune in to radio astronomy, master IRC, upgrade your monitor, follow our Pi Pico projects, and more. And don't forget to keep up to date with all our open source news and views.

## ISSUE 320 October 2024

Product code:  
**LXFDB0320**



### In the magazine

Discover how to make your home brighter and smarter with open source tech, and find out whether the AI PC is anything more than marketing speak. You can also speed up your downloads, tweak network settings, upgrade your graphics card, and emulate both a Psion PDA and the Enigma machine. We also have a *Roundup* of the best beginner distros, and so much more to enjoy.

## ISSUE 319 September 2024

Product code:  
**LXFDB0319**



### In the magazine

Get a flavour of what the latest version of Linux Mint has to offer, and dive inside Linux to discover the audio stack and learn what PipeWire has to do with it. Plus, explore the Linux filesystem, find out about NixOS, expand your storage with RAID, control multiple AI models, and record games. And we haven't even mentioned our packed reviews, Raspberry Pi and news sections...

## ISSUE 318 Summer 2024

Product code:  
**LXFDB0318**



### In the magazine

Find out how to boost your VPN privacy and set up a free VPN, and continue keeping your secrets safe with our *Roundup* of security-conscious chat services. Plus, try our tutorials on everything from tweaking your install settings to ray tracing on your Steam Deck, and plenty in between. And don't forget to read our packed news, reviews and Raspberry Pi sections to stay up to date with open source.

## ISSUE 317 August 2024

Product code:  
**LXFDB0317**



### In the magazine

Don't consign your old machines to the scrap heap – they can still be put to work with Linux. And join us as we dive deeper into Linux to explore virtualisation, containerisation and virtual filesystems. We also have tutorials on managing software and files, emulating VDUs, creating old-school pixel art, and more. Plus we have our usual news, views and in-depth reviews.

## ISSUE 316 July 2024

Product code:  
**LXFDB0316**



### In the magazine

Transform your desktop with hot KDE Plasma, and join us as we start our journey inside Linux to discover how it works. Plus, choose the best text editor for your writing projects, find out about the phone firm making an environmental impact, build your own weather machine, get back to basics with the Linux terminal, keep your video calls private, and enjoy a wealth of reviews, news and views.

To order, visit **[www.magazinesdirect.com](http://www.magazinesdirect.com)**

Select **Single Issues** from the tab menu, then select **Linux Format**.

Or call the back issues hotline on **0330 333 1113**  
or **+44 (0)330 333 1113** for overseas orders.

Quote the product code shown above and have your credit or debit card details ready.



USA? EU? THE MOON?

UK readers  
turn to  
p16

# SUBSCRIBE!

Don't wait for the latest issue to reach your local store – subscribe today and let *Linux Format* fly straight to you.



The image shows three copies of the Linux Format magazine. One is a print version, one is on a tablet, and one is on a smartphone. The cover features a penguin wearing a hard hat and holding a blueprint, standing next to a yellow brick wall. The main headline is 'FORTRESS LINUX!'. Other headlines include 'Protect your privacy with Qubes', '100 pages of Linux tricks, tips & more!', 'PLUS: HOW TO', 'AURORA RADIO', 'EASY OS CHECK', and 'GOOGLE MATTER'. A yellow circular badge on the right says '2 GREAT WAYS TO SUBSCRIBE' and 'Print + NEW digital access or digital-only!'.

» USA  
From \$135.49  
For 13 issues

» REST OF THE WORLD  
From \$135.49  
For 13 issues

» EUROPE  
From €119.49  
For 13 issues

## IT'S EASY TO SUBSCRIBE!

Visit [www.magazinesdirect.com/linux-format](http://www.magazinesdirect.com/linux-format)

Call +44 0330 333 1113

Lines open Monday-Friday, 9am-5pm GMT

Or 1-844-779-2822

Lines open Monday-Friday 8.30am-5pm EST

\*We don't actually deliver to the Moon. Yet.

## PUPPY LINUX

## Secure public Wi-Fi

**Nate Drake** helps you take your life in your hands as you navigate the murky world of public wireless hotspots.



**OUR  
EXPERT**

**Nate Drake** broke out of his cubicle at Apple eight years ago to become a freelance journalist specialising in cybersecurity and retro tech.

**A** ccording to a 2023 study by Forbes, four out of 10 people surveyed said they'd had their information compromised when using public Wi-Fi. Still, no one can deny the attraction of free, unlimited internet. The same survey found that the most common reasons people connect are to cut down on their mobile data usage.

The main risk you run when connecting to public Wi-Fi, even as a Linux user, is that your device will fall victim to a MITM (man in the middle) attack.


One common attack is to send fake ARP (Address Resolution Protocol) messages to impersonate another device, such as the network gateway. This technique, known as ARP spoofing, can be used to intercept and even manipulate your network traffic.

Many public hotspot users fancy that they're immune to these kinds of attacks if the wireless network is secured – via WPA2, for example. While this encrypts the connection between your device and the router, the Wi-Fi key for public networks is usually simple to find out. This means the attacker can be inside the same network using said encryption key, so it offers very little extra security.

This isn't to say that encryption has no value on a public hotspot. For instance, if you were to point your browser to the search engine DuckDuckGo, which is secured by TLS, snoopers could see your connection to the site but not the specific keywords you entered.

Another principal danger of public Wi-Fi is via honeypots: wireless networks set up with deliberately beguiling names, like "FREE AIRPORT WIFI", to encourage unsuspecting users to connect.

These types of rogue hotspots have become so ubiquitous that companies such as Hak5 even market devices that do this for use by pentesters. These come



Using a live distro such as Puppy Linux on public Wi-Fi ensures that no changes are saved permanently to your hard drive.

complete with web interfaces that list key statistics on devices, sites visited and so on.

This is why our first advice if you want to stay safe online is to avoid public Wi-Fi in favour of a reliable private mobile data plan. If you've no choice, though, you can reduce the risk of your Linux machine being compromised by following these simple steps.

The first step in securing Linux on public Wi-Fi is to reduce your attack surface. In other words, the fewer web services and apps that are running, the less likely an attacker is to find an exploit.

To see a list of currently running services, open the terminal and run:

```
$ sudo systemctl list-units --type=service
```

From reading the description, many of these may seem innocuous but still can present a foothold to attackers. For instance, in late September, critical flaws were discovered in CUPS (Common Unix Printing System) that could allow bad actors to remotely execute malicious code. You can disable specific services you don't need via `systemctl` – for example:

```
$ sudo systemctl stop cups
```

You can also verify that the service has been stopped – for example:

```
$ systemctl status cups
```

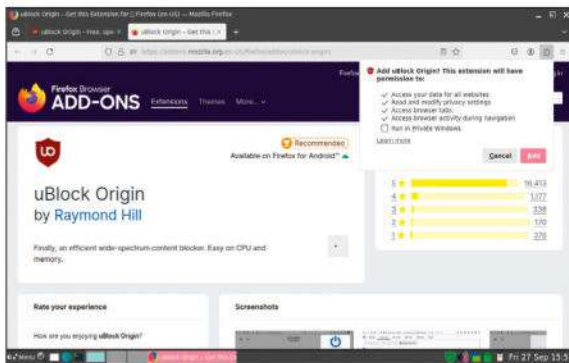
Doing this reduces your attack surface but it can be painstaking to achieve. There's also a chance that hackers will be able to find another way to remotely exploit your machine.

The best way to prevent this is to use a live version of a Linux distro when connected to public Wi-Fi. For the purposes of this guide, we've chosen a recent version of Puppy Linux (Fossapup64).

Whichever distro you choose, you need a utility such as `Etcher` to write to a USB stick. You also may

## QUICK TIP

If you're using the **Puppy Linux Terminal**, you can copy text such as commands with **Ctrl+Alt+C**. Use **Ctrl+Alt+V** or **PrtSc+Ins** to paste text. You can also just install **Gnome Terminal** via the **Puppy package manager**.



Boost your browser security with add-ons like uBlock Origin, which can automatically block both annoying ads and malicious URLs.



need to modify your UEFI/BIOS settings to allow booting from the flash drive to use the distro.

Assuming you went with Puppy Linux, it needs a few moments to copy into RAM. Although any changes you make aren't permanently saved, this works in your favour on public Wi-Fi, as if the distro is compromised, any malicious code is also lost on reboot.

As you'll see from the Quick Setup screen (see *image, left*), Puppy is also easy to secure. First tick the box marked Firewall Enabled. This blocks all incoming ports including those used by SSH, CUPS, NTP and so on. You can change this by clicking the red icon in the system tray to open firewall setup, then clicking the relevant tick boxes.

Make sure also to tick the box marked Run Internet Apps As Spot. This ensures that any web apps are run using a restricted user account. Applications run as spot can only make changes to `/home/spot`, minimising the foothold of any malware.

## Bolster your browser

The default *Pale Moon* browser should be sufficient for ordinary surfing, but you can go to Setup > Puppy Package Manager to search for and install a different browser, as we did with *Firefox*. Click Do It to install.

Upon launching the browser, you're most likely redirected to the captive Wi-Fi portal for your chosen public hotspot. This connection most likely isn't secured by TLS, so be sure not to enter any personal information unless it's required.

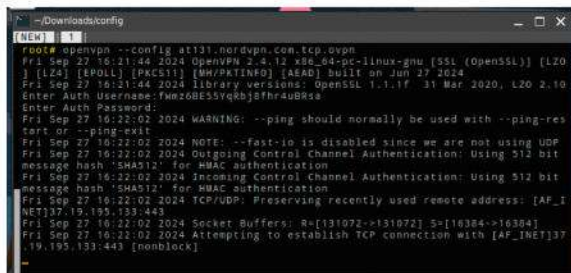
Once connected, it's time to harden your browser. You can block most annoying ads and malicious URLs with *uBlock Origin* (<https://ublockorigin.com>). If you're concerned about tracking cookies, install *Ghostery* ([www.ghostery.com](http://www.ghostery.com)), too.

## Set up your VPN

Next, it's time to secure your connection via a VPN. If set up properly, this ensures an encrypted connection between your device and the VPN server. Bad actors on the network can still intercept data packets but it is almost impossible to know which sites you're visiting or which web apps you're using.

Technically, you can just install a browser add-on for your chosen VPN. However, doing this only protects browser traffic. Plugins also don't tend to protect you from DNS leak, whereby DNS requests are processed by a local resolver instead of the VPN server.

This is why it's best to use a dedicated VPN client to encrypt all your network traffic. If you're running Puppy



```

root@pup:~# openvpn --config at131.nordvpn.com.tcp.ovpn
Fri Sep 27 16:21:46 2024 OpenVPN 2.4.15 x86_64-pc-linux-gnu [SSL (OpenSSL)] [LZO]
[LLZS] [EPOLL] [PKCS11] [MH/PTINFO] [ASAD] built on Jun 27 2024
Fri Sep 27 16:21:46 2024 library versions: OpenSSL 1.1.1f 31 Mar 2020, LZO 2.10
Enter Auth Username: 'fwm26E35Yq8bj8fhrAubKse'
Enter Auth Password:
Fri Sep 27 16:22:02 2024 WARNING: --ping should normally be used with --ping-res
tart or --ping-exit
Fri Sep 27 16:22:02 2024 NOTE: --fast-io is disabled since we are not using UDP
Fri Sep 27 16:22:02 2024 Outgoing Control Channel Authentication: Using 512 bit
message hash 'SHA512' for HMAC authentication
Fri Sep 27 16:22:02 2024 Incoming Control Channel Authentication: Using 512 bit
message hash 'SHA512' for HMAC authentication
Fri Sep 27 16:22:02 2024 TCP/UDP: Preserving recently used remote address: [AF_I
NET]137.19.195.133:443
Fri Sep 27 16:22:02 2024 Socket Buffers: R=[131072->131072] S=[16384->16384]
Fri Sep 27 16:22:02 2024 Attempting to establish TCP connection with [AF_INET]137
.19.195.133:443 [nonblock]

```

Use the OpenVPN `--config` flag to specify your configuration file, then enter your credentials to connect securely to your VPN.



Linux as per our recommendation, you can use the distro's package manager to install *OpenVPN Network Manager*.

We recommend this over using the bundled *Gpptp VPN* app as PPTP (Point-to-Point Tunnelling Protocol) has numerous security issues.

To get started, you need the necessary OpenVPN configuration files (OVPN) from your provider. Some VPN services also use separate credentials for users logging in via third-party apps.

Armed with this, launch the CLI and use `cd` to navigate to the location of your configuration file(s). You can now connect via OpenVPN – for example:

```
$ openvpn --config at131.nordvpn.com.tcp.ovpn
```

At this stage, you're asked to enter your VPN username and password.

Whichever method you use, check your VPN connection is working correctly by using a dedicated testing site such as <https://dnsleaktest.com>. **LXF**

Using a text-based browser such as ELinks will protect you from malicious JavaScript but it doesn't always work well with forms.

## »TEXT TRANQUILITY

As we've learned, the more you reduce your attack surface, the harder you'll make it for bad actors to compromise your system.

Although taking measures like hardening your browser and using a VPN can do wonders to block malicious URLs and encrypt your traffic on public Wi-Fi, you first need to get past the captive portal.

If you accidentally connect to a honeypot network, this is a particular concern, as the portal page may contain harmful scripts such as trackers to log your activity.

One way to bypass this is to use a text-based web browser such as *ELinks* (available via Puppy's package manager). Such browsers are run via the terminal so only display text and simple graphics.

The advantages of this are very clear: as no media or JavaScript loads, most tracking and exploit attempts fail.

The only downside is that some legitimate Wi-Fi networks may require you to fill in JavaScript forms in order to connect to the internet. This may or may not work depending on the website and the browser. During our tests, for instance, we were able to successfully complete a form to register a Wikipedia account, but couldn't actually complete registration because the page displayed a CAPTCHA challenge.

» **IMPROVE YOUR LINUX SKILLS** Subscribe now at <http://bit.ly/LinuxFormat>

## LIBRUM READER

Credit: <https://librumreader.com>

# Better manage your ebook library

Bookworm **Nate Drake** helps you galvanise your ebook collection and place it in the cloud with Librum Reader and Calibre.



**OUR  
EXPERT**

**Nate Drake** is a technology journalist specialising in cybersecurity and retro tech. He has taken to devouring ebooks late at night as they're not nearly as fattening as Oreos.

## QUICK TIP

Although we recommend installing via Flatpak, if you're determined to build the *Librum* client from source, you can find instructions on <https://github.com/Librum-Reader/Librum>. If you're running Ubuntu, install both `libgl1-mesa-dev` and `libxcb-cursor-dev` first.

**W**hen it comes to ebooks, in some ways Linux users are spoilt for choice. There's any number of excellent apps, such as *Bookworm* and *Foliate*, for opening EPUB and MOBI formats. The mighty *Calibre* is also incredibly useful for converting between formats, as well as downloading metadata.

Given the open source approach to OS design, however, Linux doesn't always play well with DRM (digital rights management) protected content. This may be why there's no native Amazon Kindle desktop client. Naturally, there are some workarounds. You can still use Linux versions of web browsers to manage Kindle content in your account. The desktop client can also be run through a compatibility player like *Wine* or *PlayOnLinux*, but this can cause performance issues.

In this guide, you'll discover how to download, manage and read your ebook collection with ease using our chosen program: *Librum Reader*.

## Why Librum?

Given the huge amount of software available for this purpose, we feel we should explain briefly why we've settled on *Librum Reader*.

Aside from *Librum Reader* being open source and receiving rave reviews online, during our tests *Bookworm* crashed repeatedly when performing simple operations like looking up definitions of words.

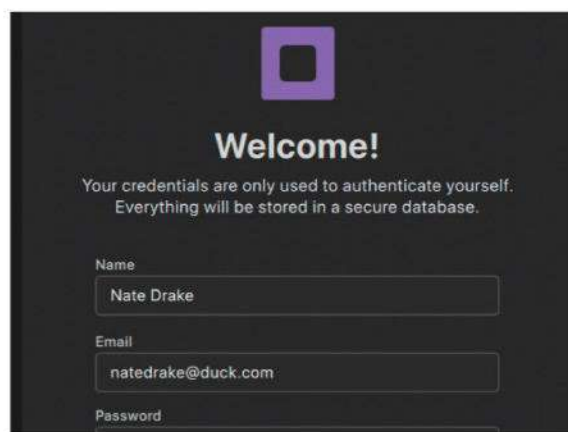
*Librum* also offers a clean, intuitive interface, making it easy to view and sort through your ebook collection. Currently supported formats include PDF, EPUB, CBZ (comic books), FB2, TIFF, MOBI and XPS.

If this weren't enough, the free pricing tier offers users 1GB of cloud storage so you can manage your ebook collection across multiple devices. Self-hosting is also supported.

## Getting started

To start organising your ebooks, head to the *Librum* main site (<https://librumreader.com>). Power users may prefer to follow the link to the GitHub page to clone the repository and compile the software themselves.

However, it's much simpler to choose Download > Flatpak. This redirects you to the *Librum Reader* page on Flathub. Click Install to download the necessary reference file. Next, minimise your browser and open



**I** Librum Reader requires a valid email address to register an account. If you don't want to share this, use an email alias or disposable address.

the terminal. Switch to your **downloads** folder with:

```
$ cd ~/Downloads
```

Double-check Flatpak is installed correctly with:

```
$ sudo apt install flatpak
```

Press Y to confirm installation if necessary. You can now install *Librum Reader* with:

```
$ flatpak install com.librumreader.librum.flatpakref
```

Enter Y to proceed. Once our install completed, we noted that the *Librum Reader* icon didn't appear in the Activities section of Ubuntu Gnome immediately. You can fix this by restarting your machine. Alternatively, run this command:

```
$ flatpak run com.librumreader.librum
```

## Add your first book

On first launch, *Librum* prompts you to log in. Click the link at the bottom to complete the registration form. You need a verified email address, but we were able to register using both an email address generated by DuckDuckGo's email protection service (<https://duckduckgo.com/email/>) as well as a temporary GuerillaMail address ([www.guerrillamail.com](http://www.guerrillamail.com)).

Once you've entered your name, email address (or alias) and your chosen password, select Let's Start. A confirmation link is sent to your email address to complete signup. Once this is clicked, *Librum* returns to the main login screen, where you can enter the



credentials you set earlier. Tick Remember Me before logging in to save having to enter this each time.

Upon login, the home screen correctly states that your ebook library is empty and prompts you to click Add Book to add titles. For now, we suggest importing only a single book from your collection.

## Edit ebook metadata

After you navigate to your chosen book and click Import, the title appears in the main window, with a cover illustration if necessary. First, select the ... at the bottom-right of the book, then Book Details. From here you can edit the metadata, such as the title and author.

You can also add supplementary information, as we did for our copy of *Alice's Adventures in Wonderland* by selecting the Document Creator as Project Gutenberg.

If the book doesn't have a cover image or you're unhappy with the current one, you can also click Change to navigate to a new image file. Choose Apply to finalise your changes to the metadata and/or cover.

Select ... once again and choose Manage Tags. From here you can add helpful tags, such as 'fiction', to the book. This is useful if you want to categorise books in a series or collection. Click Add next to each tag you type in. Tags can be removed by clicking the X by their side.

After clicking Done, check that you can search your book by tag by clicking Tags on the main screen. Any tags you entered should now be listed with a tick box so you can filter the book list accordingly.

## Customise reading preferences

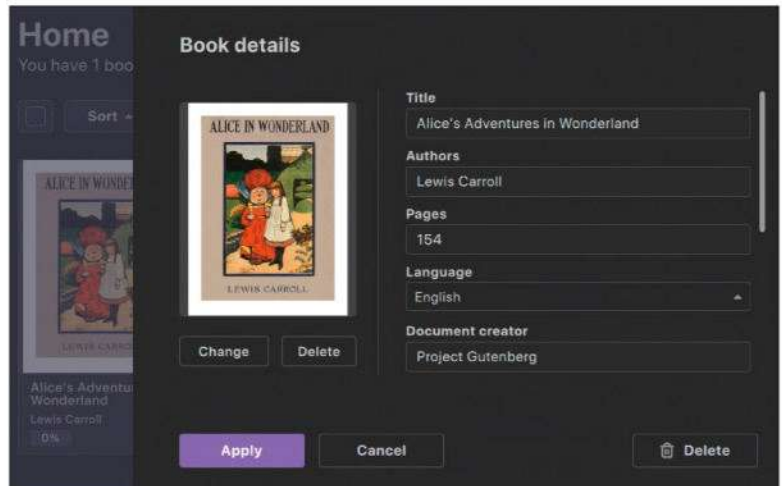
Once the metadata and tags for your first book have been tweaked to your satisfaction, double-click on the book entry itself to open it in *Librum's* viewer.

By default, the reader zooms the book text to 150%. Click the drop-down menu at the top-right as you see fit until the text is an acceptable size.

Next, click the three lines at the top-right of the viewer and choose More Options. The Display section allows you to choose between a dark and light theme.

Scroll down to the Reading section to adjust page spacing and the default zoom. From here you can also toggle displaying the book title in the title bar. The final section governs highlight colours and opacity, though we found the default settings worked perfectly.

Click the home icon in the left-hand pane to return to the main screen, then click once again into your



ebook. Take a moment to check out the icons at the top-left: the first displays a sidebar listing book chapters (if available). Click the second to view bookmarks. As this is currently empty, choose Add Bookmark at the bottom to list one of your own. By default these are named 'Bookmark' with successive numbers, such as 'Bookmark 1', but you can type a different name when setting the bookmark if you like.

Click and drag your mouse over text to see highlighting options in the available colours you selected earlier. Select a colour to highlight.

## Lookups and explanations

If you chose to highlight text in your chosen ebook in the previous section, you'll have noticed the pop-up menu also lists the ability to Look Up highlighted text.

In the case of *Alice's Adventures in Wonderland*, we put this to good effect to look up the word "hookah". *Librum* promptly displayed its definition according to Wikitionary.

This works well if you select individual words; however, if you try to select multiple lines of text or words – for example, "blue caterpillar" – *Librum* will likely be unable to find the definition.

This is where the AI Explain feature comes in useful. Users of *Librum's* free tier are allowed 10 such requests per day.

To get started, you only need to highlight a section of text and choose AI Explain. This is the default AI

Choose Book Details to edit key data such as the author and title. From here you can also choose another book cover and specify the creator.

### QUICK TIP

If you need to edit metadata for multiple books at once – to add a book description, for instance – follow the steps in the boxout (p.73) to load the files into *Calibre* first. You can then select multiple books. Right-click, then choose Edit Metadata > Edit Metadata In Bulk.

## » DRM DILEMMAS

Free and open source software such as *Librum Reader* and *Calibre* is generally only compatible with files that are free of DRM (digital rights management).

This is a divisive issue, given that if you download a book protected by DRM, you technically don't own it, so much as a limited licence to access its content.

Key players in the tech industry insist that DRM is necessary to protect intellectual property. Microsoft must have once felt this way, as in 2017 it made its own ebook store available as

part of a Windows 10 Creators Update. The store was mostly relegated to the *Microsoft Edge* browser and failed to gain any significant market share, hence the decision to close it down in 2019.

As all content was DRM-protected, Microsoft was able to unilaterally close down the libraries of all ebook customers, though it did offer refunds to affected users. The shutdown certainly went smoother than Amazon's in 2009. In an ironic twist, the tech giant remotely wiped thousands of copies of *1984* from

Kindle devices, after discovering that they'd been illegally distributed.

This Orwellian move was made possible through DRM. In the same year, Apple decided to place more trust in its users not to share content illegally by removing DRM protection from all music in its iTunes store. This is just as well these days as iTunes is no more. The same is true for other now defunct services that delivered DRM content including MSN Music, Acetrax Video on Demand and Yahoo! Music Store.



## QUICK TIP

If you've configured *Librum Reader* for self-hosting but want to change back to using the provider's servers, open the `.conf` file and change `selfHosted=false` back to `true`. You also need to change `serverHost` back to `api.librumreader.com`.

mode, where *Librum* simply describes what is happening in the text in simpler terms.

Click the Mode drop-down menu to list other AI options. For instance, choosing Explain Like I'm Five uses even more basic language to describe what's happening in a passage of text.

The Summarize AI mode deserves special mention, as it can be used to give a concise explanation of a large section of text, such as a chapter. Choose this to see the CliffsNotes version of any text.

If you change the AI mode, make sure to click Ask at the bottom-left to submit your new query.

You should also bear in mind *Librum's* admonition that "AI responses can be inaccurate". We found this to be the case when using the AI explain feature on Robert Chamber's *The King in Yellow*. The AI incorrectly claimed the protagonist was part of the Lovecraft universe.

## Organising your library

Now that you're comfortable with editing and viewing books, it's time to add the rest of your collection.

Do this by clicking Add Books at the top-right of the home screen. *Librum* actively scans books as you import them, then displays an alert if it detects duplicates. Choose Don't Add to proceed.

*Librum* doesn't support selecting multiple books at once, but you can group titles by tags. For instance, we individually edited the books *The Sign of Four*, *The Adventures of Sherlock Holmes* and *The Return of Sherlock Holmes* to share the tag 'sherlock'. You can then use the Tags menu to tick the box next to the corresponding tag to show only books in that group. Click Tags > Remove Tags to return to the default view.

By default, books are listed by the date they were added to *Librum*. This is probably easiest for identifying titles you want to read next. Still, you can change this by clicking the Sort menu at the top-left. Supported orders include percentage read, book title (A-Z) and authors (A-Z).

For the latter category, *Librum*, rather disappointingly, arranges the alphabetical list by authors' first names instead of surnames.

If you have a large number of books, you may prefer to sort them into folders. To start, click the folder icon beside the settings gear on the left-hand side.

Next click the + icon beside Folders. From here you can type your new folder name – such as 'Sherlock

Holmes'. Work your way through the Icon and Color section to make your new folder easily recognisable. You can also click Description to enter something meaningful, such as 'Sherlock Holmes Stories'.

Click Create to finalise your changes. You can now click the ... next to each book and choose Move To Folder. In the new pop-up window, make sure your chosen new folder is highlighted, then choose Select.

Click the name/icon of your new folder in the left-hand Organize pane to see all listed books. If you've added any to the folder in error, simply click the ... once again. Choose Move To Folder, then highlight ~None and choose Select to return the ebook to the main directory.

You can likewise move out of your folder into the main directory by selecting All Books from the Organize pane.

## Expand your collection

So far in this guide we've assumed that you have an existing ebook collection to import. For instance, for our tests, we downloaded a number of titles from Project Gutenberg, which contains around 70,000 public domain ebooks in various formats.

We did this the old-fashioned way by selecting files on the website for download to our test machine's hard drive. However, *Librum* has a way to access tens of thousands of free ebooks via its built-in online store.

We say store as this is what it is referred to in the documentation, but as with Project Gutenberg, it actually consists of free public domain books. To get started, click the download button (above the home icon) in the left-hand pane.

From here you can scroll through *Librum's*/Gutenberg's offerings. If you see a book you like, just click on it, then choose Download. (Note that you cannot edit the book metadata from this window.) When a book is successfully downloaded, a tick icon appears over the cover image.

Once you've finished downloading, click the settings gear in the left-hand pane, then on Storage. From here you can see your current pricing tier (Free – 1GB), as well as the number of ebooks in your collection and your remaining storage.

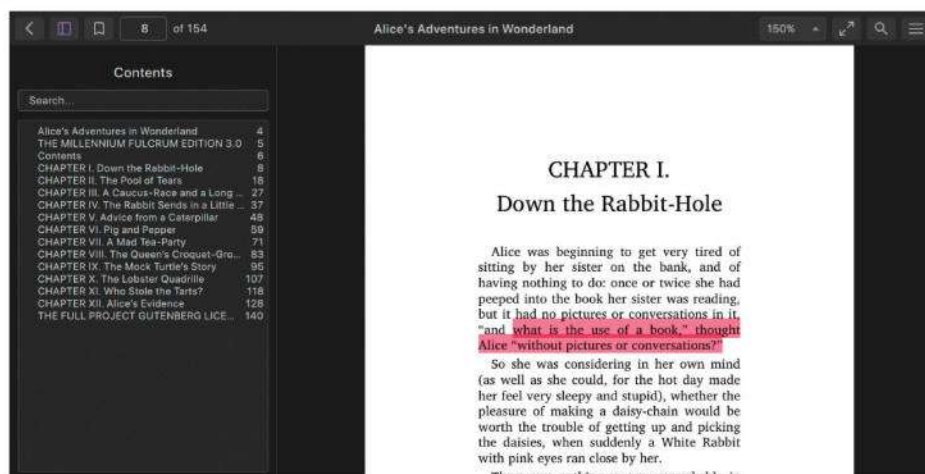
If you wish to export an ebook you've downloaded from *Librum's* store, open it in the viewer. Next click on the hamburger icon at the top-right of the screen, then the download icon. From here you can save the book to another folder on your device.

## That syncing feeling


One of the main advantages *Librum* has over simpler ebook readers such as *Bookworm* is that it can sync content to the cloud and to other devices. This should be done automatically. You can also manually sync your book collection, along with tags, progress, highlights and so on, by clicking your user icon, then choosing 'Sync'.

If you set up *Librum* on a separate Linux machine and log in using the account you created earlier, the software displays the book covers along with your reading progress. Click

Select the icon at the top-left to toggle the sidebar and display chapters, if applicable. Click and drag text to highlight in various colours.







You can also group books using folders. Select the + icon to specify a name, colour and icon. Choose All Books to return to the main screen.

on an ebook to download it, along with any bookmarks or highlighting.

If you're reading a book and want to sync your changes in real time, you can also click the hamburger icon at the top-right and then choose Sync Book.

If you prefer a self-hosted solution, the *Librum* project has also made its server software available via GitHub (<https://github.com/Librum-Reader/Librum-Server>). Manual setup instructions for Linux servers are available. However, the easiest way to get this running is via the offered *Docker* configuration file.

To get started, you need access to a suitable server. For instance, we used an Amazon EC2 instance running Ubuntu Server 22.04.

Once you're logged in to Ubuntu server, ensure you have *Docker* installed by running:

```
$ sudo snap install docker
```

Next, download the necessary *Docker* configuration file with:

```
$ wget https://github.com/Librum-Reader/Librum-Server/raw/main/docker-compose.yml
```

Create the container by running:

```
sudo docker compose up -d
```

You can now check that the *Librum* container is up by entering `sudo docker ps`.

As you'll see, *librum-server* runs on TCP ports 5000 (HTTP) and 5001 (HTTPS). It also makes use of MariaDB on TCP port 3306. If you're running an EC2 instance, as we did, make sure you edit the inbound rules via the web console to allow incoming connections to those ports.

Next, on your Linux machine close down *Librum*. You now need to edit the client configuration file to point towards the *Librum* server.

If you compiled the client yourself manually via GitHub, you'll find the necessary configuration file in `~/config/Librum-Reader/Librum.conf`.

If you installed via Flatpak, however, the file can be found in `~/var/app/com.librumreader.librum/config/Librum-Reader/`.

The easiest way to edit the file is via the terminal:

```
$ sudo nano ~/var/app/com.librumreader.librum/config/Librum-Reader/Librum.conf
```

You can also open your **home** folder and use Ctrl+H to display hidden folders to navigate there using the file manager.

Whichever method you use, be sure to change `selfHosted=false` to `true`. You also need to change `serverHost=api.librumreader.com` to the correct IP/domain name and port for your server – for example: `serverHost=3.144.191.203:5000`

If this is set up correctly, upon launching *Librum* you'll see that your library seems to be empty again (because this is a new server with a new database). You can repeat the steps above to synchronise existing content on your device.

One quirk we did notice of doing this is that if you check your available storage in Settings, it is listed as OGB. Naturally, the true amount of storage you have available depends on your server.

## Librum limitations

As comprehensive a solution as *Librum Reader* is for ebook management, it's still something of a work in progress. Aside from peculiarly labelled storage settings for self-hosting, if you go down this route you'll find that the AI Explain feature no longer works. From reading developer comments on Reddit, however, you can use your own OpenAI key to get this working.

Clicking on the link to the statistics page in the left-hand pane also reveals that it's "currently in development". It seems the *Librum* Android app is likewise a work in progress.

*Librum* does support a wide variety of ebook formats. However, if you do want to import a type that isn't supported, such as AZW3 or PDB, we suggest converting it first using a third-party tool such as *Calibre* (see boxout, below). **LXF**

## » CONVERSION CAPERS

If your ebooks are in a format that isn't supported by *Librum*, you can convert them before importing. One of the most popular tools for this is *Calibre*, which supports a huge range of file types. To get started in Debian-based distros like Ubuntu, open *Terminal* and run:

```
$ sudo apt install calibre
```

Upon first launch, *Calibre* asks you to specify your language and its working directory (by default, this is in your **Home** folder). Click Next. *Calibre* also prompts you to choose your reading device: choose Generic E-Ink > Continue, then Finish to confirm these changes.

When the main screen loads, select Add Books from the top-left and navigate to the ebook you wish to convert.

Once the book title appears in the main pane, right-click and choose Edit Metadata > Edit Metadata Individually. Double-check that both the book author and title are written correctly, and amend them as necessary.

Next select Download Metadata at the bottom of the window to have *Calibre* scour sources such as Google for a meaningful description of the book.

Return to the main window. Right-click your ebook again and choose Convert Books > Convert Individually. Double-check the output format at the top-right is compatible with *Librum* – EPUB, for example – then click OK.

Once conversion is complete, right-click the book in the main window then choose Save To Disk. The pop-up menu now shows the option to save the converted file to your hard drive.

» IF YOU'RE A BIT OF A READER... Subscribe now at <http://bit.ly/LinuxFormat>

# Upgrade it: Expand your system memory

Often found forgetting birthdays and dates, **Neil Mohr** wishes he could upgrade his own memory so easily.



**OUR  
EXPERT**

**Neil Mohr**  
first used  
Linux in an HP  
ProLiant mini  
home server  
and never  
looked back.


**O**ur upgrade tutorials are a perfect way to learn about how Linux interfaces with your PC's different subsystems. Memory is oft overlooked for understandable reasons – largely because it's taken for granted. Even if you've built your own PC, for most people it's just a case of dropping in a couple of DIMMs and letting them self-configure.

But we're hoping we can get you to stop and spare a thought for your poor unloved memory. Let's take a quick refresh of how you can interrogate memory use (both via a GUI and the terminal) and test for possible corruption that can cause random crashes. While upgrading memory in the sense of swapping out older, slower memory for faster DIMMs is an edge case, adding more DIMMs for expanding total memory can be useful. Especially in an age where 16GB of memory has become the default.

On Gnome-based desktops the *System Monitor* tool (on KDE this is *Plasma System Monitor*) offers a rich set of resource monitoring tools. Use the Process tab to see a table version similar to *top*, discussed below.

A standard Linux terminal tool to get memory details is `cat /proc/meminfo` – this returns the contents from one of the special virtual filesystem folders that actually contains hardware information. In this case, memory use on the system controlled by the kernel. A bit easier to understand is the *free* command,

A good reason for opting for four DIMM slots is that you can easily upgrade your memory down the line.



which returns more basic system memory usage. Use the standard `-help` option for more details. From the terminal, if you want an interactive display of memory use, the *top* command monitors all the running tasks (*htop* is a richer experience), alongside a host of details about memory and swap use.

Users new to Linux may be surprised to see that much of the memory is in use even on systems with low loads – this is because Linux is designed to use spare resources, and it uses any spare memory for disk buffers. When new applications request extra memory, these buffers can be quickly freed up and reallocated.


If for some reason you don't have physical access to the inside of your system, you can use *dmidecode* to list all the hardware the Linux kernel knows about and is controlling. On its own, it spits out quite a large list, but it includes motherboard and BIOS details, which can be useful in themselves. To help filter things to memory, use this (note the capitalised *Memory*):

```
$ sudo dmidecode | grep Memory
```


If you can spot lines with **Configured Memory** Speed: `xx00 MT/s`, that indicates an installed and working DIMM. An entry that shows 'unknown' indicates an empty DIMM slot or failed memory DIMM.

If your PC is suffering random crashes and you suspect it might be corrupt memory, most Linux distros offer a low-level memory test from the GRUB boot menu. To access this, reboot your PC and hold the Shift key down as it powers up. At the bottom of the various distro boot items will be Memory Test. Select that and it automatically runs, with a progress report.

If your install doesn't offer this option, there are two easy fixes. The first is to grab *Memtest86* direct (from







The classic Memtest86 can often be found in your GRUB menu.



Look, Mum, we're installing our own memory!

[www.memtest.org](http://www.memtest.org)), install it to a USB drive and boot that. Another option is if you have a Ubuntu ISO image to hand, the live disc offers the same *Memtest*.

## Upgrade considerations

No matter if you have a new-build motherboard or are considering upgrading an existing old PC (or laptop), the first key bit of information you need is the motherboard model. Do a DuckDuckGo and get the manual for the motherboard – this will tell you these key details: the type of memory it takes (DDR3/4/5), the speeds it supports (2200/2600/3200), the maximum memory size per DIMM and in total, plus the order the DIMM slots should be populated, and how many memory lanes are supported.

While all these details are important, that last point (memory lanes) indicates the optimal number to run. Typically, all motherboards support two lanes, so the processor can access two banks at the same time. Some high-end systems offer four; very old systems just one. Effectively, it means buying two DIMMs is best, but they have to be installed in the correct order to work correctly, else you'll just be getting single-lane speeds. So read the flipping manual!

Building from scratch is somewhat easier in terms of just going for the memory at the speed and price

that suit you. We've rounded up the best DDR4 and DDR5 options out there over the page. Don't pay for faster memory than your motherboard supports – unless you're planning on playing with overclocking – because you'll be wasting your money.

Upgrading an older PC gets a little more tricky but modern memory controllers are pretty flexible. If you have two existing 8GB DIMMs and another bank of two empty DIMM slots, the ideal approach is to buy two more identical 8GB DIMMs. So the same make, speed, timings, model. If you can't get the same model/make, check the timings are the same. Worst case, you can pop in any old DIMMs and ensure they're running at the speed of the slowest DIMM in the BIOS.

Modern motherboards tend to do an excellent job of configuring mismatched DIMMs. You can even run two banks of differing memory sizes – such as 2x 8GB and 2x 16GB – but Intel and AMD motherboards only provide dual-memory speeds to the lower capacity; above is single-channel. So, there is a performance hit, and again it depends on the motherboard chipset, so you need to check if it's possible at all.

If you're in any doubt, we always recommend the Crucial online advisor at [www.crucial.com/store/advisor](http://www.crucial.com/store/advisor). This is extra useful for laptops, when you can be left scratching your head. **LXF**

## » TIMINGS, TIMINGS, TIMINGS

A lot of technical knowledge is taken for granted, so let's remedy that. A memory stick is called a DIMM (Dual Inline Memory Module); reduced-sized versions are SODIMMs. The current main technology is DDR4, with DDR5 coming in for the latest Intel and AMD Zen 5. DDR3 is still very much around – it was phased out in 2015 but DIMMs can still be purchased.

On a performance level, the data rate is key – this is shown as DDR4-3600. So, the generation type followed by the transfer speed, here that

is 3,600 times per second. The motherboard chipset actually sets the speed, so no matter what speed the memory DIMM supports, it can't run faster than the chipset, and sometimes manufacturers limit this. For example, Intel H370/B360 limited speeds to DDR4-2400 and AMD Ryzen 3000 to DDR4-3600.

Unless you're a speed freak, data rate is likely all you need to worry about, but there's more! Latency is the amount of time it takes for any memory operation to

initiate, and it may come as a shock to the uninitiated that this metric hasn't changed in decades: both an ordinary stick of PC-100 and a run-of-the-mill set of DDR4-3200 have a CAS Latency of 10ns.

CAS is the time it takes to access the next bit of memory. If all rows are closed, accessing a cell requires first opening a row, and then finding the correct column (tRCD+tCAS). If the wrong row is open, accessing a memory cell requires closing the current row, opening the correct row, and finding

the correct column within the new row (tRAS+tRCD+tCAS). Finally, when command rate is increased from 1T to 2T, an extra clock cycle is required for every memory command.

A 100MHz clock cycle takes 10ns (10 nanoseconds), so that PC-100 CAS 1 took a minimum of 10ns to access data. Meanwhile, DDR4-3200 operates at a 1,600MHz clock, and a 1,600MHz clock cycle takes only 0.625ns. This means that DDR4-3200 CAS 16 takes a minimum of 16 times 0.625ns to access data, which is still 10ns.

# Upgrade it: Buying faster system memory

Zhiye Liu tests to see if DDR4 and DDR5 make a difference to your PC.

**W**hether you're shopping for the best RAM for gaming to upgrade a PC that's struggling or building a new PC from the ground up, the best RAM kit for your money depends on the platform you pick and the workloads you plan to run.

The hard part is evaluating whether faster memory improves performance. For example, if you're running an Intel system with one of the best graphics cards, most programs won't respond meaningfully to faster or slower system memory. On the other hand, some workloads scale well with higher data rates. For example, file compression programs love fast memory.

AMD's Zen CPUs benefit more from higher memory frequencies. Increased memory speeds on AMD Ryzen and Threadripper often translate to real-world gains. In games, that means higher frame rates at mainstream resolutions, such as 1080p, or smoother performance at higher resolutions. But the number of extra frames you get with faster RAM varies from game to game.

Lastly, memory speed makes a big difference if gaming with integrated graphics, whether an Intel or AMD processor. The graphics engine that's baked into most of the best CPUs for gaming doesn't generally



Ensure you check what the maximum memory limits before dropping in four DIMMs, but at 196GB you should be fine!

have dedicated memory, like discrete graphics cards do, so faster RAM also improves performance. However, if you must pay top dollar for the fastest RAM to get playable frame rates, you're better off buying slower system memory and a discrete graphics card.

So, the best RAM is usually the fastest if gaming without a dedicated graphics card on AMD Ryzen, and in some isolated scenarios with Intel. But if you don't care about squeezing out the best performance, DDR5-5600 is drop-in compatible with AMD's Zen 5 processors, and DDR5-5600 or DDR4-3200 for Intel 14th Generation Raptor Lake Refresh CPUs.

## » QUICK RAM SHOPPING TIPS

16GB continues to be the current sweet spot for many users. Programs are getting bigger and require more memory, while 1080p and 4K videos are becoming more common. PC games are also becoming more demanding, and websites are getting more complex. So, while heavy multitaskers and prosumers may need 32GB to avoid using much slower disk-based virtual memory, 16GB is far more affordable and sufficient for gaming and mainstream productivity.

Advertised XMP memory speeds might not be possible

on AMD-based motherboards. XMP is an automatic memory overclocking setting designed for Intel systems, but some mobo makers offer BIOS settings to attempt to use these faster speeds on AMD boards. However, these settings aren't present on all motherboards, and don't always work when they are. Look for memory kits with the AMD EXPO certification if you own a Ryzen 7000 (Zen 4) processor.

Dual-rank memory is faster than single-rank. Whether DDR4 or DDR5, tests show

that dual-rank kits perform better than single-rank kits on AMD and Intel platforms.

Always buy a single kit for the capacity you want. Don't combine different memory modules or kits from the same vendor and product line. Mixing and matching may not work flawlessly, and tweaking may be required to achieve stability. You'll also be opting for slower speeds to maintain compatibility.

To avoid minimum to zero manual intervention, choose a memory kit that coincides with the official memory frequency supported by your

processor. For example, DDR5-5600 is the baseline for AMD's Zen 5 chips, whereas DDR5-5600 or DDR4-3200 is for Intel's 13th Generation Raptor Lake and 14th Generation Raptor Lake Refresh processors.

Do you own a Raptor Lake or Raptor Lake Refresh CPU? Remember that these processors natively support DDR4-3200 on Gear 1. It's down to the silicon lottery if your chip can keep higher data rates on Gear 1. However, DDR5 operates in Gear 2 by default, regardless of the frequency.



# G.Skill Trident Z5 Neo RGB DDR5-6000 (2x 16GB)

When you want the best memory for your system and are happy to pay for it!



**I**t's hard to find reasons not to like the Trident Z5 Neo RGB DDR5-6000 C30. The memory kit works immaculately out of the box and offers solid performance. That alone is enough to win the majority of buyers over. Looks, as usual, are subjective, but you can't dispute the Trident Z5 Neo RGB's premium exterior. Like its competition, G.Skill uses SK Hynix M-die ICS for the memory kit, so some tweaking headroom remains in memory modules.

By default, the memory runs at DDR5-4800 with timings toned down to 40-40-40-77. G.Skill has embedded one EXPO profile into the Trident Z5 Neo RGB. Once enabled, the memory jumps to DDR5-6000, with timings and DRAM voltage fixed at 30-38-38-96 and 1.35V, respectively.

G.Skill's memory kit has some of the tightest timings on the market. As a result, the CAS Latency (CL) wouldn't move, neither did the tRAS. However, we could lower the

tRCD and tRP by two clock cycles but had to up the DRAM voltage to 1.4V. It simply wasn't worth it, so overclocking the memory kit to DDR5-6200 is a better way to squeeze out more performance.

DDR5 pricing still fluctuates, but current DDR5-6000 C30 memory kits start at £120. The Trident Z5 Neo RGB DDR5-6000 C30 has proven to be a formidable kit for AMD and Intel platforms. There will ultimately be faster kits down the road, but in the meantime, the Trident Z5 Neo RGB DDR5-6000 C30 is as good as it gets.

## VERDICT

**DEVELOPER:** G.Skill

**WEB:** [www.gskill.com](http://www.gskill.com) **PRICE:** £120

» Rating **9/10**

# Corsair Vengeance DDR5-5200 (4x 48GB)

We're going to need memory, lots and lots of memory.



**T**he typical compromise when boosting memory levels is speed; timings can't be as tight or fast when you're balancing sending bytes to four DIMMs with millisecond precision. So, while not big on speed, the Vengeance DDR5-5200 C38 is the kind of kit that will appease users who require a lot of memory, such as content creators, software developers and data scientists. Current AMD and Intel CPUs embrace DDR5-3600 and DDR5-4400, respectively, in a four-DIMM setup. Therefore, DDR5-5200 is technically overclocking, so your mileage may vary depending on your CPU and mobo.

For maximum compatibility, the modules run at DDR5-4800 with the timings toned down to 40-40-40-77. A single XMP 3.0 profile onboard will set the memory to DDR5-5200 and change the memory timings and DRAM voltage to 38-38-38-84 and 1.25V, respectively. The 38-38-38-84 memory timings are borderline stable for

this kit at DDR5-5200. However, they wouldn't cooperate with us, even with a 1.35V DRAM voltage. We also used a brute-force approach and raised the voltage to no avail.

Corsair presently sells the Vengeance DDR5-5200 C38 for £630. The company also offers a more flashy version with RGB that carries a £10 premium. It's undeniable that the memory kit costs a small fortune, but it's a valuable asset if you need lots of memory for your field of work.

## VERDICT

**DEVELOPER:** Corsair

**WEB:** [www.corsair.com](http://www.corsair.com) **PRICE:** £630

» Rating **9/10**

# Samsung DDR5-4800 (2x 16GB)



The best value DDR5 upgrade – if you can find it in stock.

**T**he Samsung DDR5-4800 C40 kit targets those who want to stick to the JEDEC baseline speed. The modules don't look pretty but have excellent overclocking headroom. There's a good potential to hit breakneck speeds with tight timings. The B-die ICs will help you push the modules. Our sample got to DDR5-5800 C36 easily, but there is no guaranteed overclock, so this kit belongs to the high-risk, high-reward category.

Samsung's memory conforms to JEDEC's baseline for DDR5. There's no need for XMP 3.0 profiles as the modules default to DDR5-4800. The timings are 40-39-39-76. The memory requires 1.1V to operate, so putting a heat spreader on it is a waste unless you're overclocking it.

For those enthusiasts that don't want to push the frequency envelope, there's also headroom for tighter timings. However, we don't see the appeal in running DDR5-4800. Nonetheless, the memory was happily operating with its timings at 34-35-35-69 when we set a 1.4V DRAM voltage in the motherboard's firmware.

The Samsung DDR5-4800 C40 modules deliver on the performance front for people looking for baseline performance, which you can install and forget. While lacking in looks, the Samsung DDR5-4800 memory has unstinted overclocking headroom. That's its most vital attribute. You can hit high frequency with tight timings.

The modules sell for around £66. That's the consumer pricing, of course, not IC pricing. As a result, a two-DIMM configuration will set you back £114 – not a bad price if you want to run JEDEC speeds. This kit can be your ticket into the B-die overclocking world.

## VERDICT

**DEVELOPER:** Samsung

**WEB:** [www.samsung.com](http://www.samsung.com) **PRICE:** £114

» **Rating 8/10**

# TeamGroup T-Force Xtreem ARGB DDR4-3600 (2x 8GB)



Best 16GB memory for AMD Ryzen builders.

**T**eamGroup did a great job with the Xtreem ARGB DDR4-3600 C14 memory kit – it certainly ticks all the right boxes. The kit looks fantastic and performs equally well when lit up or powered down. And it's the fastest DDR4-3600 C14 memory kit we've tested.

By default, the Xtreem ARGB memory modules post at DDR4-2400 with 16-16-16-39 timings. The memory modules have one XMP profile, and flipping the appropriate switch in your BIOS will get them to run at DDR4-3600 with 14-15-15-35 timings and a DRAM voltage of 1.45V.

Although the memory kits use the same ingredients, some are better binned than others. Unlike the competing

memory kits in our tests, the Xtreem ARGB could run at DDR4-3600 with much tighter timings. With a DRAM voltage of 1.46V, we got the timings as low as 13-14-14-35.

The memory market only has a handful of DDR4-3600 C14 memory kits at the 16GB (2x 8GB) capacity. And with a price tag of £59, the Xtreem ARGB is the least expensive. So, the only gripe we have is with its availability. Unfortunately, Ebay is the only retailer that lists the memory kit, so finding it could now be challenging.

## VERDICT

**DEVELOPER:** Team T-Force

**WEB:** [www.teamgroupinc.com](http://www.teamgroupinc.com) **PRICE:** £58.99

» **Rating 8/10**



# Patriot Viper Steel DDR4-3200 (2x 16GB)

The best 32GB twin DIMM DDR4 kit you'll find.



**B**uyers within the performance PC market have only two excuses to stop at DDR4-3200: either they're trying to save money or coax a good performance from a system that can't go much further.

For those in either category, but the former in particular, Patriot's 32GB Viper Steel 3200 kit outshines the competition in terms of price, beating its closest rival by approximately 18%. That's without any performance penalties beyond having the exact mid-market timings as its competitors.

Even though the Patriot Viper Steel's performance victories are less than 1% overall, its low price puts it well

ahead of even the least-expensive competitor in our basic performance-to-price comparison. As a result, value seekers within the performance PC market have just found their new champion.

## VERDICT

**DEVELOPER:** Patriot

**WEB:** [www.patriotmemory.com](http://www.patriotmemory.com) **PRICE:** £61.99

» **Rating 9/10**

# Corsair Vengeance LPX DDR4 -2666 (2x 8GB)



If you're looking for budget-friendly H370/B360 upgrades, look no further!

**I**ntel's H370 and B360 chipsets instruct its Core i5 (and above) processors to lock out any memory settings above DDR4-2666, which is particularly unfortunate in a market that's moved way past that setting. DDR4-3200 is now mainstream within the enthusiast PC market and is often treated as such by the memory sellers that cater to enthusiasts and gamers. Thus, the best way to get a top-performing brand-new DDR4-2666 kit would be to dial the way-back machine to 2016 and get the high-performing equipment from that time. Barring that, Corsair has a workaround.

Timings of 16-18-18-35 seem painfully slow compared to modern parts, but the dual-rank advantage is so great that we expect a rout with this kit. Unless you know where to find old high-end DDR4-2666, the only way we can think of to do better is to pony up for a low-latency 32GB kit, securing two ranks per DIMM and tighter timings.

Corsair's easy workaround for the performance problem of Intel's DDR4-2666 limit was to use older, low-density chips to populate its 8GB DIMMs with two ranks rather than the single rank of its competitors. The best part is that they did this without a significant price increase. We recommend it for anyone whose XMP-compatible platform has a maximum DDR4-2666 data rate, which includes most retail boxed H370 and B360 motherboards. **LXF**

## VERDICT

**DEVELOPER:** Corsair

**WEB:** [www.corsair.com](http://www.corsair.com) **PRICE:** £35.99

» **Rating 8/10**

## Manage RAID services on your Linux servers

A sometimes overlooked skill, **Stuart Burns** explains the basics of preparing and creating pools of RAID storage.

**M**ost sysadmins use hardware RAID controllers, which are quick and easy. That's fine but if a RAID controller goes south, it can be difficult to get a replacement that recognises the proprietary RAID format the controller uses. Hardware RAID controllers also increase the purchase costs.

Software RAID used to create a noticeable overhead. With the advances in CPU tech and tuning, for most workloads the overhead is now unnoticeable.

Linux has software RAID built into the kernel via the *mdadm* tool. The beauty of Linux soft-RAID is that the RAID-enabled disks can be read in any Linux system you can plug the disk set into without encountering any RAID hardware dependencies. An administrator can even configure a Linux soft-RAID environment as boot disks. As an example, it's in Ubuntu server installations under Something Else on the disk configuration screen if you don't want to do it manually.

We're covering the basics of implementing Linux soft-RAID using a pair of disks for non-boot RAID (because we don't have enough space here).

One of the things a lot of new Linux administrators get wrong is that they create the RAID from the raw disks rather than using partitions.

Using partitions rather than raw disks means the administrator can set the size of all the partitions to the size of the smallest disk. In future, if you replace a failed disk, it's as easy as creating an equally sized partition

on the replacement and resynchronising. No nasty 'force'-style commands to make the disk fit.

It may sound trivial but even disks from the same manufacturer of the same size and type can differ by a small amount and create headaches.

### Prepare the disks

Identify the disks using `sudo lsblk`. Ours are **sdc** and **sdd**. It is critically important to ensure the right disks are identified. The device in question (assuming it's a new empty disk) shouldn't have any branches on it, like disk **sdf** in the first screenshot (*opposite page, top*).

On each disk, create partitions of the same size using the *parted* command. Here, we are using *parted* to create a partition of 10TB on each disk:

```
$ parted /dev/sdx # Replace x with the appropriate disk
mklabel gpt
(parted) mkpart primary 0GB 10TB
quit
```

Repeat the same on the other disk, substituting **sdx** for the second disk. Check the partitions are correctly reporting using `sudo lsblk` again. This time you should see something like the second screenshot (*opposite page, bottom*), with **sdc1** and **sdc2** with two identically sized partitions.

### Create the mirror

Post partition setup, we can use *mdadm* to do the heavy lifting of creating the software RAID. Here the command to create a simple disk mirror is shown:

```
$ sudo mdadm --create /dev/md1 --level=mirror --raid-devices=2 /dev/sdc1 /dev/sdd1
```

To demystify this key command a bit, the **create** parameter specifies the new virtual disk to be forged from the union of the two disks. By default, a randomly selected value for **md(random number)** is used if not explicitly specified.

The **--level** parameter is key. It specifies the RAID level to be used. It can be striped (don't use this unless you understand the implications – it creates one big disk but it also increases the risk of failure and is not resilient at all). We are using mirrored – two identical disks to provide resilience against hardware failure. Detailed information on RAID levels and their use can be found at <https://bit.ly/47GeXP2>.

The last parameter specifies the number of drives and the partitions to make up our desired drive layout.

### » YOU'RE A BAD BLOCK!

It is always a good idea to thoroughly test your disks before using them. This is especially true with disks of unknown provenance. Linux has a command called *badblocks* that can be used as a basic fitness test to make sure there are no bad blocks on the disk.

As the name implies *badblocks* can be used to give the disk an extensive workout and test for failures by writing patterns on disks and re-reading

them, across the whole disk, potentially multiple times.

Note: this is a destructive command – it destroys all the data on the disk in the process:

```
$ sudo apt install -y badblocks
$ badblocks -w -s /dev/sdX
# Substitute for your drive here.
```

There is a non-destructive *badblocks* test that can be run by emitting the **-w** parameter. It is not as effective without the write, but it can still be useful to test for bad blocks.

Warning: this process can take a while. Creating a mirror comprising two 10TB drives took 24 hours. This is to be expected and doesn't mean the newly created disk isn't performant. It's just the up-front setup process. (While it is being built, it can be mounted and used.)

The status of the build (and the RAID sets in general) can be found in the file `/proc/mdstat`. Just echo out the file to see how the build process is going:

```
$ cat /proc/mdstat
```

A bit of a pro tip: prefacing the above command with `watch` shows a constant refresh of the build progress (see image below).

On the same subject, notice the `[UU]` in the image. It looks innocuous but each `U` represents a disk. An upper-case `U` means it's healthy. An `_` means a disk is dead, dying or missing.

Once the build is completed we can build a filesystem on it:

```
$ sudo mkdir /data
$ sudo mkfs.ext4 /dev/mapper/md1
$ sudo mount /dev/mapper/md1 /data
```

If so desired, you can manually add it to the `fstab` for automatic mounting on bootup. Just treat it like any other disk, so in `/etc/fstab` it would look similar to this:

```
$ /dev/mapper/md1 /data defaults 0 0
```

To prevent a potentially non-booting scenario, check the `fstab` is correct by using `mount` to check everything mounts correctly:

```
$ sudo mount -a
```

Assuming no issues are thrown up, everything is good to go and will persist post-boot.

## Replace failed pool members

It's all very well having a RAID setup, but a device will eventually fail. To replace it is quite easy and consists of the following steps. To get more information on the failure, use the command:

```
$ sudo mdadm --detail /dev/mdx # Replace mdx with the failed disk
```

Remove the failed disk from the array. Add in the replacement disk and create a partition as we did earlier. Use the following command to rebuild:

```
$ sudo mdadm --manage /dev/md0 --add /dev/sda1
```

Lastly, let's talk about monitoring. By default, there is no automated reporting (other than potentially syslog, depending on the configuration). Due to running out of space here, we will come back to this and cover it next month. For those who want the bare details, you can run a cronjob to verify the status and send an alert when it detects an issue. **LXF**

```
sysadmin@media:~$ sudo lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINTS
sda	8:0	0	119.2G	0	disk	
└─sda1	8:1	0	118G	0	part	
└─md0	9:0	0	117.9G	0	raid1	
└─vg0-bootlv--0	253:2	0	110G	0	lvm	/
└─vg0-lv--0	253:3	0	4G	0	lvm	[SWAP]
sdb	8:16	0	223.6G	0	disk	
└─sdb1	8:17	0	1M	0	part	
└─sdb2	8:18	0	118G	0	part	
└─md0	9:0	0	117.9G	0	raid1	
└─vg0-bootlv--0					lvm	/
└─vg0-lv--0					lvm	[SWAP]
sdc					disk	
└─sdc1					part	
└─md127					raid1	
└─data					crypt	/data
sdd	8:48	0	9.1T	0	disk	
└─sdd1	8:49	0	9.1T	0	part	
└─md127	9:127	0	9.1T	0	raid1	
└─data	253:4	0	9.1T	0	crypt	/data
sde	8:64	0	9.1T	0	disk	
└─sde1	8:65	0	9.1T	0	part	
└─backup	253:5	0	9.1T	0	crypt	/backup
sdf	8:80	0	3.6T	0	disk	
nvme0n1	259:0	0	931.5G	0	disk	
└─nvme0n1p1	259:1	0	931.5G	0	part	
└─trashdisk-temp_vol	253:0	0	50G	0	lvm	/tmp

Use `lsblk` to identify the correct disk. It should usually be the one with no 'tree' hanging from it.



**Stuart Burns** is a sysadmin for a Fortune 500 enterprise based in that London.

## » LICENCE TO KILL

As I write this, the blogging universe is having a torrid time after two big names in the WordPress world (WP Engine and WordPress.org) have had a very major, very public fallout. The story boils down to (among other things) WP Engine refusing to contribute to the cost of maintaining the WordPress infrastructure that enables all of WP Engine's plugins and so on to be updated in the WordPress environment, despite being financially able to.

It has now gone as far as **WordPress.org** banning WP Engine updates being served from its infrastructure, causing many issues, not least the fact that security updates can't be directly delivered to end users. The thing is, **WordPress.org** is unlikely to be hurting for cash.

Yes, it's free software, but it isn't immune to the fallout and recriminations of two already very rich companies with large founder egos in charge. It's not a good look, but the big takeaway is that even in the world of open source software, understanding which external vendors your software uses is critical.

Just because it is open source and often even free as in beer doesn't mean it doesn't come without risks. Other examples of risks include that, of late, a lot of companies that used to be open source have now mutated into some form of rather 'smashed together' open source but not free-to-use-as-you-wish licensing. *Terraform* versus *OpenTofu* is one example. Embrace and change the licence to make more cash seems to be the order of the day.

Red Hat is another example, if one were needed. It's enough to turn people on to BSD licensing!

```
root@media:~# cat /proc/mdstat
Personalities : [raid1] [linear] [multipath] [raid0] [raid6] [raid5] [raid4] [raid10]
md1 : active raid1 sdd1[1] sdc1[0]
      9766302720 blocks super 1.2 [2/2] [UU]
      [====>.....] resync = 18.1% (1774575168/9766302720) finish=1450.8min speed=9
1803K/sec
      bitmap: 62/73 pages [248KB], 65536KB chunk

md0 : active raid1 sdb2[1] sda1[2]
      123665408 blocks super 1.2 [2/2] [UU]
      bitmap: 1/1 pages [4KB], 65536KB chunk
```

Building the array. Get used to this screen, because it takes a while to complete!




# tom's **HARDWARE**

**GET ALL THE ESSENTIAL BREAKING NEWS FOR THE TECH ENTHUSIAST!**

No matter if you're building a PC, buying a laptop or learning about robots, Tom's Hardware has all the comprehensive knowledge you need.



**Scan & Subscribe  
for free!**



## **GET READY TO BECOME A TRUE PYTHON EXPERT**

Make Python work for you with tutorials on coding with Django, Flask, Pygame and even more useful third-party frameworks.

**ON SALE  
NOW**

Ordering is easy. Go online at:  
**magazinesdirect.com**  
Or get it from selected supermarkets & newsagents

# HotPicks

PDFSam » CPU Info » Qrcp » Recordbox »  
Newsboat » Safe Eyes » TLP » Hurry Curry! »  
HyperPlay » Geeqie » Rclone



## Mayank Sharma

has been a closeted hoarder for decades, but instead of junk in his house, he stuffs his machines with apps and utilities.

### MANIPULATE PDF

## PDFSam

Version: 5.2.8

Web: <https://pdfsam.org>

For those who frequently work with a large number of PDF files, there's a number of operations that we need to perform on a regular basis. This includes rotating and extracting pages from a PDF file, splitting a large file into smaller chunks, or merging several PDF files into a single one.

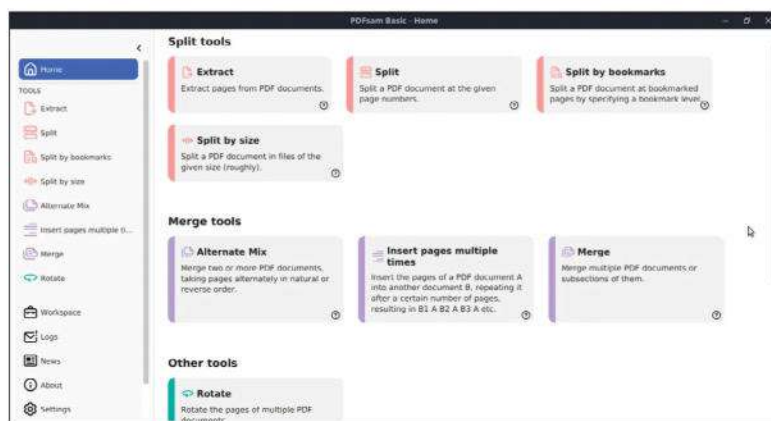
PDFSam is a fast and featureful utility that offers these functionalities in an easy-to-use and intuitive interface. The project only supports 64-bit machines, and needs at least 256MB RAM to operate – although we know from experience that having more RAM is always a plus for merge or split operations.

Many popular desktop distributions carry PDFSam in their repositories, but Ubuntu and Debian users will appreciate the availability of the DEB package, which can be installed with the `sudo dpkg -i <pkg-name>` command. For all others, the project offers a portable binary. Download the portable archive (.tar.gz) from the project's website, and uncompress the files. Navigate into the `pdfsam-5.2.8-linux/bin/` directory and you can now launch PDFSam simply by running the `sh pdfsam.sh` command.

When you launch PDFSam, you're greeted by the Dashboard, which features large icons that describe all the features available to use. The app offers various options for its split and merge features. For instance, you can opt to split a PDF file by size, page numbers or bookmarks, or even extract specific pages from a PDF. There are a lot of sane and easily discernible options for each feature, and the helpful tooltip on mouseover helps make sense of each option.

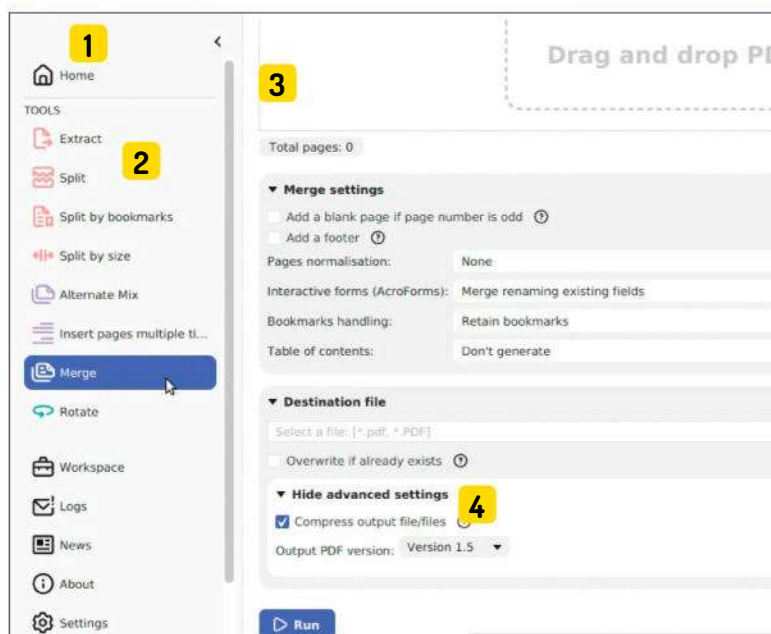
After you select a feature in the sidebar or Dashboard, PDFSam offers additional feature-specific options, such as the destination filename. Press Run at the bottom of the interface to let the app perform the specific operation.

PDFSam Enhanced is a premium variant that offers extras such as OCR, signing documents, creating PDFs from DOC, PPT, XLS and other formats, and more.



PDFSam is incredibly useful and speedy for basic PDF editing operations such as extract, split, merge and rotate.

### LET'S EXPLORE PDFSAM...



#### 1 Sidebar

The sidebar provides quick access to all the available features. Click Settings at the bottom to tweak app behaviour, or Home at the top to return to the Dashboard.

#### 2 Available features

All available features are listed in the sidebar under the Tools heading. The project's website provides a quick and short explanation of each.

#### 3 Feature-specific config

When you select a feature, you can drag and drop files at the top, and tweak the settings. The available settings vary depending on the feature.

#### 4 Advanced options

Users will appreciate the option to compress PDFs under the advanced settings. This feature can be used to reduce the size of large PDF files.



## SYSTEM MONITOR

## CPU Info

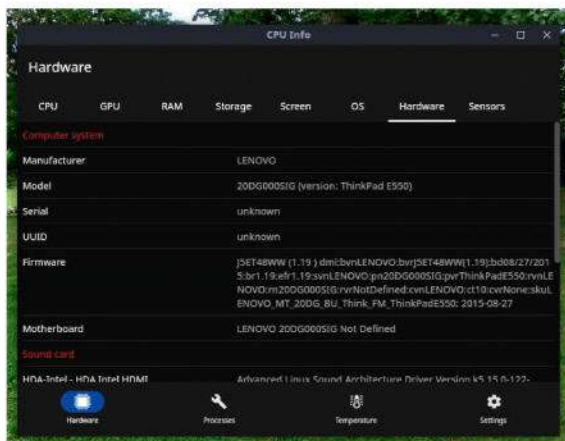
Version: 1.3.0 Web: <https://github.com/kamgurgul/cpu-info>

You won't find this system monitor in the software repositories, but thankfully installation is straightforward with the Flatpak package. Run the `flatpak install flathub com.kgurgul.cpuinfo` command first, and then `flatpak run com.kgurgul.cpuinfo` to launch it.

When you launch *CPU Info*, you're presented with information about the hardware, and the interface features a series of tabs at the top such as CPU, GPU, RAM, Storage, OS, Hardware and so on. You can click on each to find pertinent details. For instance, the RAM tab lists the total and available memory. You can find details about the manufacturer, model, firmware and motherboard, as well as info about the soundcard, network interface and so on under the Hardware tab.

Not all parts of the app work as advertised, however, as *CPU Info* failed to identify the orientation of the Ubuntu 20.04 LTS test machine, and reports the OS to be FreeDesktop SDK and manufacturer as GNU/Linux.

The interface also features a series of buttons at the bottom. The first is Hardware, which is what *CPU*



Click the Settings button on the bottom-right of the interface to define the temperature unit and switch to the dark theme, as we've done here.

*Info* defaults to. The Processes button lists only two: flatpak-bwrap and java. One assumes these are both *CPU Info*-related, as we're running the app through Flatpak, and it's written in Java. The Temperature button at the bottom shows the CPU temp, which you can also access by clicking the Hardware button, then the Sensors tab on the top-right. Surprisingly, *CPU Info* reports different temperatures at both these places.

The project has no documentation to speak of. If you need help or assistance, head over to the Issues page on GitHub and read through the resolved problems. If needed, you can also open a new issue, and hopefully you'll find answers. As it stands, the project offers little value or substance over other such utilities, which are far more featureful.

## FILE TRANSFER

## Qrcp

Version: 0.11.3 Web: <https://github.com/claودیodangelis/qrcp>

Unlike the old days when a single laptop or desktop formed the entirety of one's work environment, modern workspaces span multiple devices and operating systems. One has to constantly share files and content across these devices. As useful as USB cables are for transferring files, they don't offer a quick solution when you have to share files from a Linux desktop to an iPad, for instance. What if you could share files across devices, irrespective of the underlying operating system, with nothing more than a QR code? This is the functionality *Qrcp* provides.

You need Go 1.18 or higher on your system to install *Qrcp* with the `go install github.com/claودیodangelis/qrcp@latest` command. If that seems like too much work, you'll appreciate the DEB and RPM packages released for ARM and 32- and 64-bit machines.

The source and destination devices both need to be on the same Wi-Fi network for *Qrcp* to work as it binds a web server to the machine's IP address, along with a random port number. *Qrcp* then generates a



unique QR code that provides the relevant information. You can then read the QR code on your mobile device, which automatically takes you to the decoded URL. The web server automatically stops once the download is completed.

To transfer files from the desktop, run the `qrcp <filename>` command. You can also provide a space-separated list of files to transfer with the same command. Next, scan the generated QR code on your mobile device, and you're taken to the web page from where you can download the specified file or a zipped archive, if you opted to transfer multiple files.

You can similarly run the `qrcp receive` command on the desktop to receive files from a mobile device.

*Qrcp* is one of the simplest file transfer utilities we've come across, and can be used to send and receive files without hassle.



## MUSIC PLAYER

## Recordbox

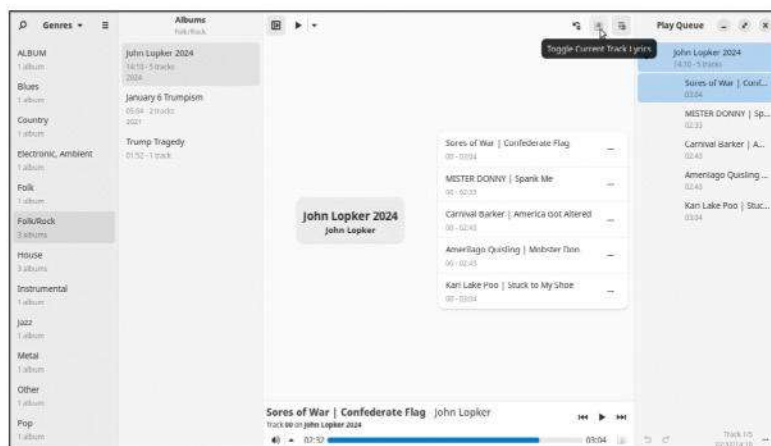
Version: 0.8.1 Web: <https://codeberg.org/edestcroix/Recordbox>

Just as good music is a matter of taste, and one that's quite subjective, music players are designed to cater to different needs as well. Unlike most other desktop music players, *Recordbox* is designed with playing albums in mind, rather than tracks or playlists.

In a stark departure from most other desktop apps, the only way to install *Recordbox* is via the Flatpak package. You won't find the app in the software repository of your desktop distro, and the project doesn't provide binaries to build it from source.

From the terminal, run the `flatpak install flathub ca.edestcroix.Recordbox` command to install *Recordbox* and then run the `flatpak run ca.edestcroix.Recordbox` command to launch it.

When you first run the app, you're asked to point it to the music library or directory where you store your music files. The app then automatically organises the albums. The drop-down at the top of the left-most pane enables you to switch between the Genre or Artist layout. All albums in the selected Genre or Artist



are listed in the next pane. The third pane is the player, and the fourth, enabled by default, is the player queue.

You'll want to run *Recordbox* in full-screen mode, to which it defaults, in order to access all the panes.

Click the hamburger icon on the left-most pane, and select Keyboard Shortcuts for a list of all configured shortcuts, which are split into different categories, such as Navigation, Play Queue, Playback and so on. Hit Ctrl+, to open the Preferences dialog box.

One of the best features of *Recordbox* is that it enables you to edit the lyrics that a song is playing, if the lyrics are part of the metadata. You can alternatively save lyrics to the metadata from within the app itself. Look for the Toggle Current Track Lyrics button at the top-right of the player pane.

While most music players essentially have a similar look and feel, the album-centric approach of *Recordbox* lends it a different layout.

## RSS FEED READER

## Newsboat

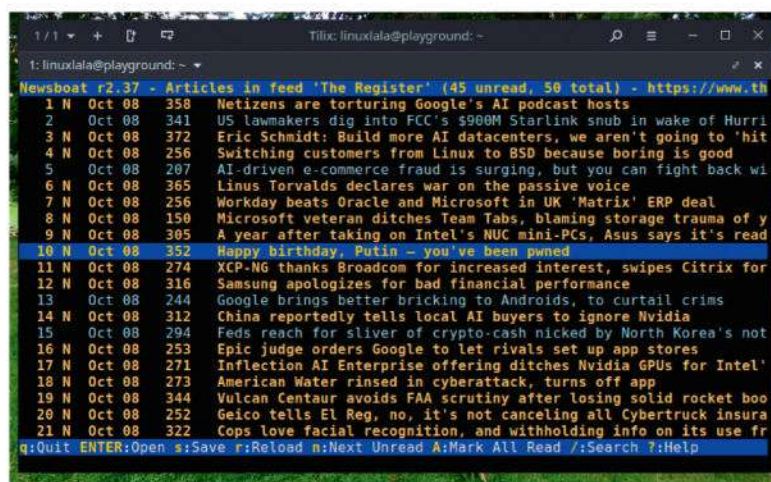
Version: 2.37

Web: <https://newsboat.org>

There are two types of internet users: those who get their information from podcasts, and those who prefer the written word. *Newsboat* is an RSS/Atom feed reader for the terminal. Forked from *Newsbeuter*, the MIT-licensed project is available in the software repositories of most Linux distros, such as Fedora, Ubuntu, Arch, Debian and others.

If your distro doesn't carry the latest version, you can compile it manually, but you have to first satisfy all its dependencies, such as *make*, *Rust*, *Cargo*, *SQLite3*, *libcurl*, *pkg-config* and so on. Refer to the project's website for the complete list of dependencies. If you don't want to go through all that trouble, you can take advantage of the Snap package and run the `sudo snap install newsboat` command to install *Newsboat*.

The default feed location varies depending on how you opted to install *Newsboat*. It's either `~/newsboat/urls` or `~/snap/newsboat/xxx/newsboat/urls`. You can also create the `urls` file at any other location of your choice and then point *Newsboat* to it with the `newsboat -u <feeds-file>` command.



All feeds must be on a separate line, and you can use tags to easily identify the feeds. When you now run *Newsboat*, it presents a list of all available feeds. Press Shift+r to refresh the stories for all the feeds, or select a feed from the list using the arrow keys, then press **r** to refresh the stories for just that feed. Hit Enter to navigate into the feed and access all the articles.

The app lets you search through feeds for articles of interest and you can even filter out specific content.

The default settings should be sufficient for most users, but you can create your own config file with relevant customisations if you prefer. Refer to the documentation on *Newsboat*'s website for information about configurable parameters.

Some websites prevent feed readers from downloading the entire article, so you can open the article in the default browser by pressing 'o'.

## EYE-STRAIN PREVENTION

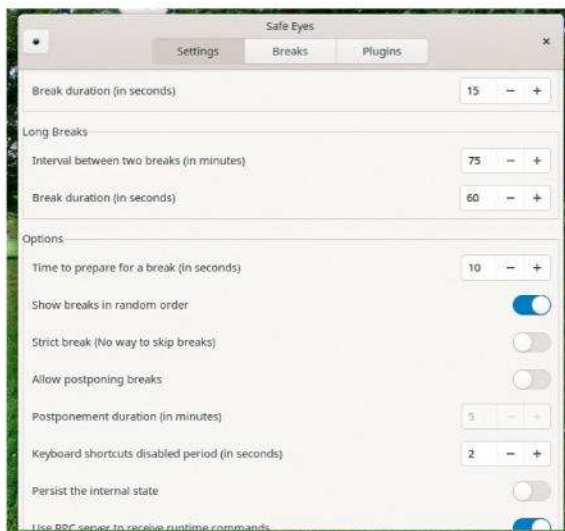
# Safe Eyes

Version: 2.2.1 Web: <http://slgobinath.github.io/SafeEyes/>

**T**oo much of a good thing is not an idiom one often associates with computing devices, but it fits. If you spend too much time staring at the computer screen, you're likely to develop a number of physiological issues, such as poor posture, weak glutes, bad core strength and, worse still (*conspiracy theories?—ED*) you'll strain your eyes. Thankfully, there's an easy fix, at least so far as your eyes are concerned.

*Safe Eyes* is a nifty utility that helps prevent eye strain by periodically forcing you to take a break. It differs from mere reminder utilities because *Safe Eyes* blacks out and locks your screen, forcing you to rest your eyes, or at least not continue staring at the screen.

The Python utility can be found in the repositories of most popular distros, so if you already have *Pip3* installed, you can run the `sudo pip3 install safeeyes` command to get the latest version. There's a Flatpak package on offer as well, which you can install with the `flatpak install flathub io.github.slgobinath.SafeEyes` command, but the project reports that many plugins and features don't work well with the Flatpak variant.



The Plugins tab has even more useful features, such as Do Not Disturb, which disables *Safe Eyes* if the active window is in full-screen mode.

Once installed, *Safe Eyes* starts automatically whenever you boot into the system. Click on the system tray icon and Settings to tweak *Safe Eyes* to your needs. The settings are split into three tabs: Settings, Breaks and Plugins.

By default, the tool forces you to take a short break every 15 minutes, and long breaks every 75 minutes. You can easily change this from the Settings tab of the *Safe Eyes* config window. You can also define the duration of the breaks, in seconds, and choose whether or not you can skip the forced breaks. Additional break-specific settings can be accessed from the Break tab.

## BATTERY OPTIMISER

# TLP


Version: 1.7.0

Web: <https://linrunner.de/tlp>

**T**LP is a battery optimising utility that you only think you don't need. However, you'll quickly change your opinion when you see how much it improves the battery life on your notebooks and laptops.

You'll find *TLP* in the software repositories of most popular desktop distros, such as Ubuntu, Fedora, OpenSUSE and so on. There are distro-specific installation instructions on the project's website, which essentially boil down to using the distro's package management utility to install *TLP*. If your distro doesn't offer the latest iteration, you can fetch the source tarball from the project's GitHub page (<https://github.com/linrunner/TLP/releases>), extract the files with `tar xzvf TLP-1.7.0.tar.gz` and run the `sudo make install` and `sudo make install-man` commands to install it.

You can now start *TLP* with the `tlp start` command. Depending on your distro, the utility might require that you manually enable the service so that *TLP* starts automatically on each reboot. To do that, run the `sudo systemctl enable tlp.service` command.



The configuration is stored in the `/etc/tlp.conf` file, and we highly recommend you don't make any changes with a 'what does this do?' philosophy on a production machine.

*TLP* works by tweaking kernel parameters that affect power consumption. As such, there's a high likelihood of compromising your system if you start tweaking settings without understanding them. Thankfully, the default settings should suffice for most desktop users. These include enabling settings or profiles based on certain events, such as when you plug in power, suspend the system, or insert a USB device, and so on. You can enable or disable events-based optimisation with `tlp-rdw disable`.

For a list of all configurations, along with information about disks, CPU, graphics card and so on, you can run the `tlp-stat` command. Refer to the man page for a quick introduction to the different categories of information available and the command switches you can use to restrict the output, such as `tlp-stat -b` to view battery information.



**RESTAURANT SIM**

# Hurry Curry!

Version: 2.1.1 Web: <https://hurrycurry.metamuffin.org>

**W**e've come across many games where the objective is to build cities or empires, but never one that involved preparing meals. If you know who Gordon Ramsay and Nigella Lawson are, you're the target audience for *Hurry Curry!*

Released under the AGPLv3 licence, the quickest way to install the game is with the Flatpak package. Run the `flatpak install flathub org.metamuffin.hurrycurry.client` command to install the game, and the `flatpak run org.metamuffin.hurrycurry.client` command to run it.

At first launch, you have to sign the employment contract, which provides the hairstyle you must adhere to, your position, and other pertinent details, such as duties and compensation. It probably won't surprise you to learn that your duty is to serve customers the meal or item they request, and your compensation is \$0.00 per month. Sign the contract at the bottom-right and you're ready to begin.

The game uses WASD keys for movement, while the arrow keys are used to pan the camera around.



You can change these and learn the other controls by heading into Settings > Controls.

Click Play and connect with the online public server, or press Start Server on the bottom-left and then connect to the *Hurry Curry!* server.

Customers enter the restaurant and start ordering food as soon as you begin. Your assignment is to prepare each meal by cutting, cooking, baking and combining the resources in the kitchen.

The project's website is home to a 10-page recipe book for such delicious treats such as tomato soup, burger, curry with rice, and so on, although on our first venture with the game, the customers all ended up asking for pizzas and burgers. You lose points if you take too long to prepare the meals.

We're embarrassed to admit that we managed to hit -5 points in about the time it took to grab this screenshot.

**GAME STORE**

# HyperPlay

Version: 0.20.0

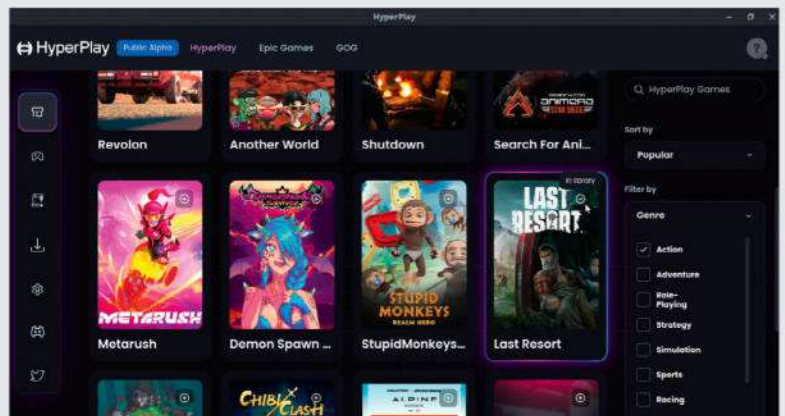
Web: [www.hyperplay.xyz](http://www.hyperplay.xyz)

**G**aming on Linux has gained tremendous traction in recent years, with many excellent treats now on offer. *HyperPlay* is a game store and launcher with a vast number of games on offer. It also supports the entire library of both the Epic Store and GOG.

The cross-platform game store is still in alpha, so you won't find it in the software repositories of most popular desktop distros. But that's quite all right as the project publishes a Snap package, and installable binaries for RPM and DEB-based distros, but recommends installation using Flatpak.

You can install it with the `flatpak install flathub xyz.hyperplay.HyperPlay` command, then the `flatpak run xyz.hyperplay.HyperPlay` command can be used to launch the store.

When you start *HyperPlay*, the first step after choosing the language and deciding whether to let the application send pertinent usage data to the server to improve functionality is to set up a wallet to store your digital assets such as NFTs and tokens, but you can



skip this step by clicking the Skip For Now button at the bottom.

You'll want to run the app in full-screen mode to fully appreciate its offerings. The pane at the right of the screen lets you choose games based on genre, such as action, adventure, strategy and so on, and sort the selection by name or popularity.

There's no way to get *HyperPlay* to list games available for the Linux platform, but when you find a game you like, click Add To Library, then read the specs, such as download size and storage space required. If you decide to continue, hit Install. Games are installed in the `~/Games/` directory by default. Hit the Play button to launch the installed game.

*HyperPlay* automatically launches Wine if the chosen game is Windows-based.

You need to have an Epic or GOG account if you wish to access those games via *HyperPlay*.



## IMAGE VIEWER

## Geeqie

Version: 2.5


Web: [www.geeqie.org](http://www.geeqie.org)

With professional-looking photos being taken on a near-daily basis thanks to ever-improving smartphone cameras, you need a featureful image viewer and organiser to make sense of all your clicks. Most Linux distros offer several image viewers in their repos, and you might find one installed out of the box. *Geeqie* is a lightweight image viewer and organiser that prioritises speed and image previewing.

Released under the GPLv2 licence, the project started in 1998 as *GqView* but adopted the new name in 2007. It's been under constant development since and you'll appreciate the thorough documentation on offer that discusses its myriad features.

Although *Geeqie* is available in the software repositories of most distros, the version on offer might not always be the latest. While you can build it from source manually, the project also provides *Applmage* as well as Flatpak packages for swift installation.

Run the `flatpak install flathub org.geeqie.Geeqie` command followed by `flatpak run org.geeqie.Geeqie` to work with the Flatpak package.



The default interface features a sidebar on the left, with a menu bar and toolbar atop it. The sidebar features a folders pane under the toolbar, which can be used to navigate to the directory where you store images. The lower part of the sidebar lists all the images in the selected folder, while the main section of the interface on the right is the image pane.

You can zoom in or out, or set the image to fit the screen using the buttons on the toolbar. These features, and more, can also be accessed by right-clicking on the image in the image pane, which also shows the assigned keyboard shortcut for each.

We've used *Gwenview* a lot, and are stunned with the performance advantage of *Geeqie* in comparison.

*Geeqie* also lets you change the layout, and offers four options. Click **Edit > Configure This Window**, and select a layout that better suits your workflow.

## REMOTE BACKUP

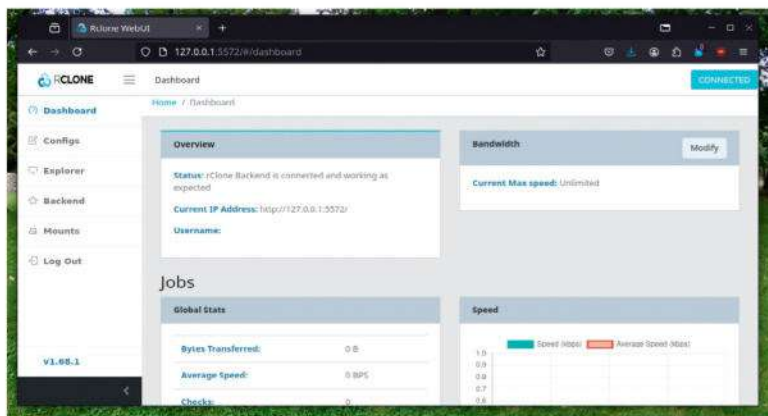
## Rclone

Version: 1.68.1

Web: <https://rclone.org>

There's no such thing as being too cautious when it comes to safeguarding data. Although a lot of people stress over using secure passwords and even encryption, not many bother with keeping secure copies of data. Unfortunately, disk troubles don't announce themselves in advance, and you can't anticipate a hastily executed `dd` command, which could deprive you of your precious data, to say nothing of malicious attacks and external threats. Backups are the only sane solution to ensure your data is protected from all threats, known or otherwise.

*Rclone* is a nifty utility that supports dozens of cloud storage services to help you safeguard your important data on remote sites. Most distros carry *Rclone* in the software repositories, but you can also grab the latest tarball from the project's GitHub page, extract the file, and build it manually. To make the installation process simpler still, *Rclone* also provides an installation script. Run the `curl https://rclone.org/install.sh | sudo bash` command to install the latest version. There are even RPM and DEB packages on offer to accommodate all





Linux users, irrespective of their familiarity with Linux distros.

Run the `rclone config` command to start the configuration process.

The configuration step provides a list of all the supported cloud services where you can host your backups. These include all the popular ones you're probably already familiar with, such as Amazon, Microsoft OneDrive, Google Drive, Dropbox, Nextcloud, OVH and many more. You're guided through the process of connecting to the cloud service of your choice, with quick access to the helpful documentation where required.

You can mount each backup as a disk, mirror the cloud data to local disk, and even use tools such as *ncdu* to analyse the data. **LXF**

You can run the `rclone rcd --rc-web-gui` command if you'd rather work with the browser-based GUI.



**Can you help inspire the  
next generation of coders?**



**Code Club** is a nationwide network of volunteer-led after school clubs for children aged 9-11.

We're always looking for people with coding skills to volunteer to run a club at their local primary school, library or community centre for an hour a week.

You can team up with colleagues, a teacher will be there to support you and we provide all the materials you'll need to help get children excited about digital making.

There are loads of ways to get involved!  
So to find out more, join us at **[www.codeclub.org.uk](http://www.codeclub.org.uk)**



# Completely complete your shell commands

With his heart bursting from his chest **Ferenc Deák** reassures us it's not a medical emergency but being so close to finally completing this series!



### OUR EXPERT

**Ferenc Deák** feels this project is his Sistine Chapel, only with less of the Pope shouting at him and more Mancunian editors.

**C**ommand line completion is like having a butler in the terminal, ready to anticipate your every command and whisk it into action with a few keystrokes. It's a bit like magic.

Type a snippet of a command or file path, press Tab, and voilà! The rest appears as if summoned by a wizard's incantation. Or by a few lines of code hidden in plain sight, like we will present in this iteration of our shell saga.

Command line completion comes in two flavours:

**Tab completion:** The old faithful. Just tap the Tab key, and our shell does its best to complete your command or file path faster than you can say, "Abracadabra." If there's more than one possibility, just stress the Tab key a little bit more, and it circles through all the available commands or directories.

**Option completion:** Think of this as the fancier cousin of Tab completion. Tap Tab twice very quickly after you have entered a command, and a menu appears, like a delightful spread at a dinner party, offering a selection of completion options upon which to feast your eyes.

### Classical Tab completion

The easiest and most widely supported by current shells of these two different completion mechanisms is the classical Tab completion, which enables users to complete the current command line with a single press of the Tab key with the most appropriate command that can be guessed from the context of the typed-in sequence. However, even in this simple scenario, we have to consider several alternatives:

1. The user typed in a partial sequence of a command, which is matched upon a list of the commands that are available at that certain location. This location, however, must also include all the directories in the **PATH** environment variable.
2. The user typed in a directory name, then pressed Tab after it. This would give an indication to our shell that it needs to enumerate all the files in that

```
fld@soleo-lin ~-> apt
autoremove (Remove packages no longer needed as dependencies)
changelog (Download and display package changelog)
depends (List package dependencies)
edit-sources (Edit sources list)
...and 11 more rows
```

The Fish shell showing off its skills with its own options autocomplete, complete with descriptions.

specific directory and offer them within the Tab completion options.

3. The user typed in a partial command, preceded by a relative or absolute directory. In this case, we need to identify all the possible completions in the directory, and cycle through them accordingly.

### Partial commands

To handle this, we need to gather all the executable files that can be found in the directories pointed to by the **PATH** environment variable. The simplest way to achieve this is with the following C++ code:

```
const char* path_env = std::getenv("PATH");
std::vector<std::string> path_dirs =
    splitStringWithDelimiter(path_env, ":");
for (const auto& dir : path_dirs) {
    for (const auto& entry : std::filesystem::directory_
        iterator(dir)) {
        auto executable = entry.is_regular_file() &&
            ((std::filesystem::status(entry).permissions() &
                std::filesystem::perms::owner_exec) ==
                std::filesystem::perms::owner_exec);
        if (executable) executablesInPath.push_back(entry.
            path()); } }
```

The code uses the **std::getenv** method to retrieve the value of the **PATH** environment variable, then uses **splitStringWithDelimiter** to split it by the delimiter ":" then traverses each directory with a **directory\_iterator** (not the **recursive\_directory\_iterator** we used in the plugin retrieval phase). The interesting piece of code is the one determining whether a file is an executable, and if yes, we add the file to our vector of executable programs found in all the **PATH** directories.

### QUICK TIP

Discover the full project and much more at its GitHub site: <https://github.com/fritzone/lxf-shell>



Once this is up and running, we can proceed to the command reading, which we heavily modified in the previous instalment of this series.

The Tab key has the key-code 9, so we need to add specific implementation for this, too:

```
if(keys.size() == 1) {
    if(keys[0] == 9) {
        command = resolveShortcut(command, "Tab"); } }
```

The **resolveShortcut** method is the reincarnation of the **resolveHistory** method from our previous episode. Now it also handles the Tab completion requests, not only the history. It has been extended to handle the Tab shortcut, which in itself is not a difficult task, and it runs along the lines of the following function:

```
std::string tabCodeCompletionResult(const
std::string& command) {
    static std::string currentCommandSequence =
    command;
    static int currentVectorIndex = -1;
    static std::vector<std::string> result;
    If(command != currentCommandSequence ||
    currentVectorIndex == -1) {
        currentCommandSequence = command;
        result.clear();
        currentVectorIndex = -1;
        for (const auto& element : executablesInPath)
            if (element.find(command) != std::string::npos)
                result.push_back(element); }
    currentVectorIndex++;
    if(currentVectorIndex < result.size())
        return result[currentVectorIndex];
    else {
        currentVectorIndex = -1;
        if(currentVectorIndex < result.size())
            return result[currentVectorIndex]; }
    return "";
```

The function **tabCodeCompletionResult** provides command completion suggestions by searching a predefined list of executable commands (the vector **executablesInPath** is the one we built in the previous step) for matches to the given command string. It uses static variables to track the current command sequence and the index of the next suggestion. If the input command changes or the function is called for the first time, it regenerates the list of matches. The function then increments the index to return the next suggestion from the list. If no more suggestions are available, it resets the index and starts over.

The whole idea behind Tab command completion can be explained with this very short method. The other two possibilities above are just a variation of this scenario, so for brevity, and to not cause confusion, we will keep the current implementation as is, but if you are in the mood, and have the free time and resources, you are more than welcome to implement the required functionality and push the code to the project's GitHub site: <https://github.com/fritzone/lxf-shell>.

## Command line options

The second completion framework is a different kind of beast. Modern shells all offer the possibility to inspect the command line options of various well-known applications. For example, the screenshots illustrate how Fish and Zsh handle the pressing of the Tab key after typing in the magical three letters 'apt'

```
> $ apt
autoclean      dist-upgrade  list          remove        update
autoremove     download     policy        search         upgrade
build-dep      edit-sources  purge         show
changelog      full-upgrade  rdepends       showsrc
depends         install      reinstall     source
```

(the installer for Debian and relatives). And yes, there is a space after the 'apt' that activates the magic.

The developers of these shells have dedicated hundreds of man hours to provide the autocompletion scripts for all the possible combinations of available Linux commands, but unfortunately we have neither the time nor resources to do this. So, we will just present a possible solution and framework for this feature and we sincerely hope there is someone out there who has the time, resources and determination to implement this for all the more than thousand different Linux commands.

Our solution depends on creating the option completion files for each significant Linux command by gathering these options in a file and letting our shell go through it while identifying potential targets for the code completion. The option completion file is just a huge JSON file, created specifically for this purpose.

## JSON structure

Below is an explanation of the JSON format:

**Root object:** The JSON structure begins with a root object ({}). Inside this root object, there is a single key called "commands".

The Zsh shell showing off its options autocomplete for the apt command in this instance.

## » STATE MACHINES

A state machine, also known as a finite state machine (FSM), is a computational model used to design systems with a limited number of specific states. It consists of a set of states, transitions between these states, and actions that result from these transitions. State machines are particularly useful for modelling the behaviour of systems that respond to a sequence of inputs or events, ensuring that the system is in a well-defined state at any given time.

In a state machine, each state represents a specific condition or situation in the system. Transitions between states are triggered by events or conditions, and can have associated actions that are executed as the transition occurs. This makes state machines an effective tool for designing and analysing the behaviour of complex systems such as control systems, interfaces and protocol design.

State machines can be represented visually using state diagrams, where states are depicted as nodes, and transitions as directed edges between them. They can also be described mathematically using tables or formal languages. Overall, state machines provide a clear and systematic way to manage the states and behaviours of a system, ensuring predictability and reliability in its operation.

For our shell, a simple implementation of a state machine can be found in the handling of the key input, where depending on which key is pressed, the shell goes into different stages, such as handling Tab completion, acting upon history, appending to the current command, or the final state, represented by Enter, where the key input state machine stops and sends the input to the next stage of the shell, command execution. As can be seen in the code, the state machine mostly consists of a series of **if** statements, but more complex state machines have more complex class hierarchies, each instantiation representing one state from the machine's finite states.



## » JSON'S FAMILY

While JSON (JavaScript Object Notation) is a widely used, lightweight and text-based data interchange format known for its simplicity and ease of use, several alternatives exist, each offering unique strengths for different use cases. Here's an overview of these alternatives.

**INI:** A simple, human-readable format for config files, though it lacks support for complex data structures. In C++, the `iniparser` library is commonly used to work with INI files. While INI is easy to use, it is outdated compared to YAML or TOML.

**XML:** XML is a flexible format that supports complex data structures and is often used in web services and configurations. In C++, libraries such as `libxml2`, `TinyXML2` and `RapidXML` are commonly used to parse and manipulate XML data. XML is more verbose and complex than JSON but offers robust features such as schema validation.

**YAML:** YAML is highly readable and is often used for configuration files. For C++, `yaml-cpp` is the primary library for parsing and emitting YAML data. YAML's human-friendly syntax makes it a popular choice for configuration, although it can be error-prone due to its indentation sensitivity.

**Protocol Buffers (Protobuf):** Developed by Google, Protobuf offers strong typing, schema support and efficient serialisation. Protobuf is ideal for complex, large-scale systems where performance and schema evolution are critical, though it is less human-readable than JSON.

**Avro:** Avro is a binary serialisation format that also supports schema evolution, making it suitable for big data environments. The `avro-cpp` library is used in C++ to manage Avro data. Avro is efficient for large datasets but is less human-readable than JSON.

**TOML:** TOML is a minimal, human-readable configuration file format that is more structured than INI. In C++, the `tomlc99` library can be used to parse and write TOML files. TOML is favoured for its simplicity and ease of use in config files.

**CBOR (Concise Binary Object Representation):** CBOR is a binary data format that is both compact and efficient, designed for small code size. The `tinycbor` library is used in C++ to handle CBOR data. It is used in environments where performance and compactness are essential, although it is not human-readable.

**BSON (Binary JSON):** BSON extends JSON with additional data types and is primarily used in environments such as *MongoDB*. In C++, `libbson` is the go-to library for working with BSON data. BSON is more efficient than JSON for certain operations but is slightly larger in size.

Each of these formats has its own strengths, and the associated C++ libraries enable developers to integrate these formats effectively into their projects, whether the focus is on human readability, performance or support for complex data structures.

**Commands array:** The value associated with the **"commands"** key is an array (`[]`) that holds multiple command objects. Each object in this array represents a command line utility with its details.

**Command object:** Each command object contains the following keys:

**"command":** A string representing the name of the command, such as **"ls"**.

**"executable":** A string representing the full path to the executable for the command, such as **"/usr/bin/ls"**.

**"arguments":** An array of argument objects, where each object represents an argument that can be passed to the command.

**Arguments array:** The **"arguments"** key holds an array of argument objects. Each argument object contains the following keys:

**"argument":** A string representing the argument itself, such as **"-a"** or **"--all"**.

**"type":** A string indicating the type of the argument, such as **"short"** (for example, **-a**) or **"long"** (**--all**).

**"description":** A string that describes what the argument does.

**"options" (optional):** An array of option objects, which only appears if the argument has specific options associated with it.

**Options array:** Oddly this is optional. Some arguments have an **"options"** key, which contains an array of option objects. Each option object contains:

**"option":** A string representing the option itself, such as **"always"**, **"auto"** or **"never"**.

**"description":** A string that describes what the option does.

### Let's breakdown an example

Here's a breakdown of a specific part of the JSON that is the beginning of the command line completion file:

```
{
  "commands": [
    {
      "command": "ls",
      "executable": "/usr/bin/ls",
      "arguments": [
        {
          "argument": "-a",
          "type": "short",
          "description": "Do not ignore entries starting with ."
        },
        {
          "argument": "--color",
          "type": "long",
          "description": "Colorize the output",
          "options": [
            {
              "option": "always",
              "description": "Always colorize the output."
            },
            {
              "option": "auto",
```

```

    "description": "Colorize the output only when
the output is a terminal."
  },
  {
    "option": "never",
    "description": "Never colorize the output."
  }
]
}
]
}
]
}
```

Let's take you through what each line is doing here:

**Command:** "command": "ls": The command is *ls*.

**“executable”:** `“/usr/bin/lis”`: The path to the executable is `/usr/bin/lis`.

**Arguments:** The `ls` command has two arguments here:

**First argument:** “argument”: “-a”: The argument -a is a short form.

**"type": "short"**: Indicates it's a short form argument.

**“description”: “Do not ignore entries starting with.”:**  
Describes the argument’s effect.

**Second argument:** “argument”: “--color”: The argument **--color** is a long form.

**"type": "long"**: Indicates it's a long form argument.

**“description”:** **“Colorize the output”:** Describes the argument’s effect.

**Options for --color:** “option”: “always”: One possible option for --color.

**“description”:** “Always colorize the output.”: Describes this option.

Upon application startup, this JSON file is loaded into memory, and a vector of C++ structure hierarchies is constructed from it (**commandsForCompletions**), each entry representing a command, an argument and a specific option for the argument. The following code details this (**clc** stands for command line completion):

```

class clc_option {
public:
    std::string option;
    std::string description;
    clc_option(const std::string& opt, const std::string&
desc)
        : option(opt), description(desc) {}
};

class clc_argument {
public:
    std::string argument;
    std::string type;
    std::string description;
    std::vector<clc_option> options;
    clc_argument(const std::string& arg, const
std::string& t, const std::string& desc)
        : argument(arg), type(t), description(desc) {}
    void add_option(const clc_option& opt) {
        options.push_back(opt);
    }
    std::string format() const
    {
        return ANSI_COLOR_WHITE + argument +
            ANSI_COLOR_NO_UNDERLINE + " " +

```

```

        ANSI_COLOR_DKGREY + description + ANSI_
COLOR_RESET;
    }
};

class clc_command {
public:
    std::string command;
    std::string executable;
    std::vector<clc_argument> arguments;
    clc_command(const std::string& cmd,const
std::string& exec)
        : command(cmd),executable(exec) {}
    void add_argument(const clc_argument& arg) {
        arguments.push_back(arg);
    }
};

std::vector<clc_command> commandsForCompletions;

```

When all this falls into place, we just need to hijack the Tab key (again) and see whether the command we typed in is the **commandsForCompletions** vector, and if yes, just print out the arguments of the command.

As mentioned, the only thing that remains for the completion of the autocompletion feature is a full-time employee, who takes all the several hundred default commands on a default Linux installation, then takes the default arguments for all the default commands, and completes the default autocompletion JSON file with them, thus giving our shell the much-needed push forward in user-friendliness.

On a standard Linux distribution with a typical desktop environment and a reasonable set of tools installed, there are around 1,500 to 2,500 commands available. This number increases significantly if you install additional software, specialised tools or programming environments.

For the moment, this is the stage at which we'll leave command completion; maybe someone will finish it with interesting features, such as option completion for the current argument, or validation if someone types a string instead of a number.

Or maybe they could make it nicer, like how Zsh handles it, by presenting a list and allowing the user to select the appropriate completion item, because there is more than one user-friendlier implementation.

## Almost complete...

Next time, we will embark on an exciting theoretical journey of exploration, focusing on how to introduce new features that will bring our shell closer to becoming a fully functional and practical tool. We'll delve into the key aspects that not only enhance its usability, but also provide a more powerful and user-friendly experience. These enhancements will transform our shell from a basic command interpreter into a versatile tool that can be adapted to various workflows and use cases. By the end of the final instalment, you will have the requirements for a shell that not only meets the needs of everyday users but also offers advanced capabilities for power users and developers. And then you can jump on a keyboard and implement all that we didn't have the time or resources to include during the lengthy span of this series. **LXF**



# Test if you're human with our PHP captcha

David Bolton hopes to capture your imagination as he explains how to create a PHP captcha that can help to cut down form spams.



## OUR EXPERT

**David Bolton** used to create PHP websites a long time ago. Luckily, they let him off for good (programming) behaviour and he promised never to create another one again.

## QUICK TIP

If you are passing important parameters between pages, always use Post not Get methods, and include a method of obscuring them, such as hashing, so they can't be easily changed.

**C**reating a captcha is something of a challenge; you need to balance it being too difficult for a computer program to solve against not being so hard that a human can't pass its test.

Most of us have probably experienced difficult captchas that feature bizarre-looking fonts that take two or three attempts to solve. Having both upper- and lower-case letters that are almost indistinguishable is particularly evil. That breaks factor two.

The captcha in this tutorial was developed because an online contact page was getting lots of spam messages. Almost all spam comes from scripts, so anything that can slow them down is going to help the issue. The big problem is coming up with something that is easy to solve for a human but hard for a computer program, and we thought that asking a simple question might work.

So, our captcha is going to output a question in text form, but what sort of question should it be? It has to be relatively easy to answer, which rules out trivia questions. Any trivia we could think of would be too specific; writing a trivia question that anyone can answer is quite difficult. It might, for instance, involve British royalty, but as visitors to the website could come from anywhere in the world, they might know nothing about that. So, let's forget trivia.

After some thought, we decided on simple mathematical operations, such as `4 + 5 = ?`. But in that simple text form, it would be a little too easy. Instead, let's put it the numbers as text, like this:

**What is four plus five?**

To solve that requires something that can extract the text, make sense of the question and then come up with the correct answer. Possibly the likes of *ChatGPT* and so on might be able to do that, but this is only meant to be an example; it's a speed bump, not a complete road block.

## Let's make it a bit harder

What if the text is output as a graphic, rather than text? Normal text on a web page is output typically between `<p>` and `</p>` tags, like this, as text characters:

```
<p> What is four plus five?</p>
```

## Captcha example

four plus seven equals?

Answer in text:  E.g. thirty two

A PHP captcha in action.

This is very easy for a web-scraping bot to read. However, if the question is output as a graphic instead of text, then the web scraper needs to be able to do optical character recognition (OCR) to extract the text from the graphic. That's something us humans can do naturally but it takes a lot more effort for a web-scraper tool.

And to make it a little harder, we can draw lines through the text to try to obscure it somewhat. So, we end up with something that looks like the screenshot (above) on the web page.

## Implementing the captcha

The code for the captcha is in a file called `captcha.html` which is a mixture of HTML and PHP. By default, your setup might not be able to execute this as PHP, but it can be configured to do so by configuring Apache web server. There are various methods how to do this described on the Stack Overflow page at <https://bit.ly/lxf322stack>.

The captcha is on the author's website, so you can try it yourself at: <https://dmbolton.com/ula/captcha.html> which calls `check.php`.

If you look at the HTML with View Source in your browser, you get something like this, minus the usual HTML boiler plate text:

```
<form method="post" id="form" name="form"
action="/ula/check.php">Answer in text:
<input type="text" placeholder="Enter answer"
name="answer" id="answer">E.g. thirty two<br>
<input type="hidden" value="<?php echo
$c->hashed($c->convertNumberToWord($c-
>answer));?>" name="correct" >
```

```
<input type="submit" id="submit" name="submit">
</form>
```

Earlier in this file it includes **g.php**, which is an include file with the code that generates the image.

## Displaying text as an image

This is the code that generates an image in **g.php**. It uses image functions supplied from PHP 4 upwards. It's in a function **image(\$txt)** where **txt** is our question:

```
$img = imagecreate(500,100);
$textbgcolor = imagecolorallocate($img,255,255,255);
$textcolor = imagecolorallocate($img,0,0,0);
$font = './Montserrat-Regular.ttf';

imagefttext($img,20,0,5,90,$textcolor,$font,$txt);
imageline($img,0,80,500,80,$textcolor);
imageline($img,0,75,500,75,$textcolor);
imageline($img,0,85,500,85,$textcolor);
ob_start();
imagepng($img);

printf('',base64_encode(ob_
get_clean()));
imagedestroy($img);
```

The font used came from **www.1001freefonts.com**. It must be uploaded to the same location as the PHP files used or else you'll have to specify the path in the line that starts **\$font =**.

The **imagefttext()** call draws the text, which is then overwritten by the three **imageline()** statements. These draw horizontal lines over the text. Feel free to draw diagonal lines or shapes – in fact, anything you like to partially obscure the text and make OCRing it harder. The **ob\_start** and **ob\_get\_clean** functions capture the output into an output buffer and then pass it into the **printf**, which outputs the image.

## Object orientate this bad boy

All of the code to generate the captcha has been put into a PHP class called **Captcha** in **g.php**.

Before the image is generated, these lines create an instance of the **Captcha** class and pass an array of strings called **\$tests** into the captcha constructor, where it gets stored in an internal array called **\$captchas**.

```
$c= new Captcha($tests);
$txt=$c->parse();
$c->image($txt);
```

The array **\$tests** is defined:

```
$tests = array("a+b", "a-b", "a*b", "d+b", "d-b", "re");
It's referred to in code by $this->captchas[$index].
```

## Picking a random sum

The function **Captcha.pick()** returns a random element from **\$this->captchas[]**. If you want to add in more types of sums, just add them in the array **\$tests**. The more the merrier.

The mathematical numbers are mostly of the form number operation number – like **a+b** – but some don't have an operator, like **re**. The number symbols have a meaning: the **a** symbol is replaced by a random number in the range 1-15, but if the operation is a subtraction, then **b** (a random number in the range 1-10) has to be less than **a**. If **b** is greater than **a** and the operator is a subtraction, it adds another random number to **a**, so **a-b** is always positive, so no minus sign needed.

The symbol **d** is short for a dozen and **e** is a random number in the range 1-9 that is then squared. The symbol **r** means root, as in square root. So, the question might be what is the square root of 16, and the answer is four.

The function **\$captcha->parse()** does these things:

1. It gives the symbols a value according to the rules.
  2. It then builds up the question as text in the variable **\$result**.
  3. It calculates the answer. That is needed for the second page to check the answer.
- The only complications are with the **b** symbol, which always happens after the operation
- If we were to add another sum – for example, half – we might do it by adding **ah** into **\$tests**. This takes

### QUICK TIP

Especially when dealing with strings and hash functions, use **echo**, **print** and **print\_r** for debugging your code. An irritating bug turned out to be a space that was added on the end of a number after converting to text. This badly affected the hashed value.

## » WHY IS IT SO HARD TO WRITE A CAPTCHA?

Ever since programmers started writing web bots, captchas have existed in one form or another. Google originally used difficult-to-read text from books and learned how to decipher it from people's choices. In the last few years, the company has gone for a selection of photos that include something to pick out in some of them, such as a tree, bus, crossing and so on. In theory, this is difficult for a computer program but easy for a human.

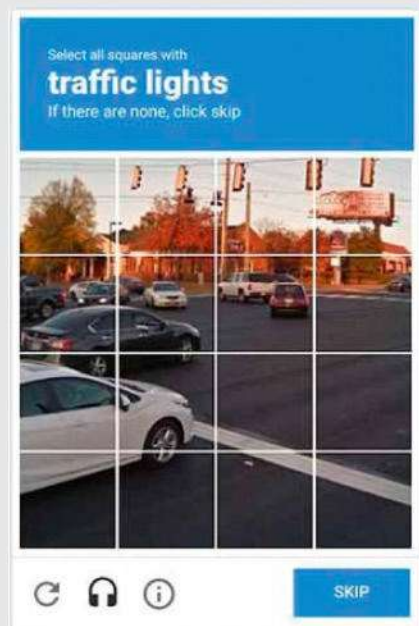
However, AI has made it easier for software to be more human-like. The difficulty is creating a challenge that defeats a program but isn't too hard for a human.

Plus, you have to realise that this is an arms race and no matter how good your

captcha is now, someone will defeat it eventually. Currently, we favour an approach that means the captcha-solving bot has to extract text, read it and then perform an action. In this particular case, the action is solve a simple mathematical operation. To make it more challenging, the mathematical operation is picked randomly from a list each time and the numbers are also generated randomly.

We don't expect this to be foolproof, but if you use it, it will keep bots out for a while. It's not going to be the best captcha in the world, but our example just shows you how to develop one.

A typical Captcha, where you need to identify which images contain a specific object.



App icons, most of which are written in PHP.

## » HOW WE STOPPED WORRYING AND LEARNED TO LOVE PHP



### QUICK TIP

Seeing a blank page during development? Use the `ini_set` for `display_errors` and `display_startup_errors`, and `error_reporting(E_ALL)` to help identify bugs during development. You can remove them or set it to a different level once you are happy with your code.

PHP is a back-end scripting language and is available with all Linux hosting. This author started with PHP about 20 years ago, beginning with PHP 4, and wrote quite a few sites. However, back then, PHP had its issues.

The language had been poorly designed (compare it to Python, which was superbly designed), it wasn't typed, and had lots of inconsistencies in function naming. PHP got a lot of flack over that. Read <https://eev.ee/blog/2012/04/09/php-a-fractal-of-bad-design/> from 12 years ago

to see how bad it was. To get things done, you had to spend a lot of time going through the online documentation and looking for examples in the comments.

In spite of that, PHP made possible many excellent open source packages, of which WordPress is probably the best known.

But starting with version 7, things improved considerably, including the speed. PHP was never that slow, but 7 and then 8 made it faster still. The OOP became more consistent and this

has continued with version 8. Improvements included named arguments, union types, attributes, constructor property promotion, match expression, null-safe operator, JIT, and in the type system, error handling and consistency. Version 8 was released almost four years ago and it's now up to 8.3.

This captcha was the first PHP program this author had written in 15 years and it worked more or less first time. There was none of the hassles that we remember from before.

the value `a` and then divides it in two. We don't want to deal with any pesky decimal points, so if `a` is an odd number (detected by the test `a%2==1`), we increment `a`. The `$result` text would be something like "What is half of eight".

To add this in, under the first `switch` statement in `parse()`, add the following block under the `break` for `case 'r':`

```
case 'h':
if ($left%2==1) $left++;
$result = "what is half of " + $left;
$right="N/A";
break;
```

Because there is no operator needed here, the variable `$operator` is not changed.

If you add any sums that use `b`, make sure it is always after the operator.

Also, to increase the difficulty a little bit more, you can play around with the text of the questions – for instance, you can use these different ways of saying the same thing; just randomly pick one of the following phrase structures:

```
What is eight plus five
Calculate eight plus five
Eight plus five equals
The sum of eight and five is
```

Basically, anything that makes it more difficult for the computer program to understand the text is good for your captcha.

### Making a hash


We have to pass both the correct answer and the typed-in answer through to the next page, so we do this by passing the correct answer as a hidden variable. The entered answer is sent in the `Post` statement. Anyone could view the source and see the value of the hidden variable, so we need to disguise it so they can't. The answer value is hashed by calling `Captcha->hashed($answer)`. To stop anyone trying out various hash functions, a salt string "XYZ1234" is appended to the answer in word form and passed into the hash function, which uses `ripemd256` as the parameter `$correct`.

The first parameter passed into the `hash()` function is the name of the hash algorithm used. To see what is available in your PHP installation, run a bit of PHP code that calls `print_r(hash_algos())`:

```
<?php
print_r(hash_algos());
?>
```

This outputs a list of all of the hash functions available. Avoid using `md5` and `sha1` as they were





Our enhanced captcha throws in shapes and colours – apologies if you're colour-blind.

### QUICK TIP

If all else fails with your testing, take a look in the folder (on the PHP server) where your code is and look for the `error_log` file. This is a text file with all your sins listed. If you can't see the file, use the `phpinfo()` function to return information about your server.

cracked nearly 20 years ago. The comments on the bottom of this hash function documentation page show typical outputs for each hash function. For **ripemd256**, it's a 64-character value – generally, the longer the value the better, but 64 characters is more than enough.

Hashing is a one-way function that converts a string such as **fiveXYZ1234** into a meaningless 64-character string of text. Feel free to use your own salt (the name for the text **XYZ1234**), but make sure you change it on both PHP scripts.

There's not much to say about the function **convertNumberTo Word(\$num)**. It just converts a number like 5 into text like five. So long as a numeric value is provided, it converts that number into one or more text words. If there are two words in the text, like 32 when converted to thirty two, they are space-separated. The answer must have a space between when there are two numbers.

On the next page (**check.php**), the values are read from the `$_POST[]` and if the hashed value of the answer plus salt matches the **\$correct** value, it's good. **Check.php** also includes **g.php** to get access to the **hashed** function.

### Checking the answer

The two values **answer** (what you entered) and **correct** – the hashed value of the correct answer – are passed using the post method.

After retrieving the values in **check.php**, this line does the test to see whether it's correct:

```
if ($c->hashed($answer)== $correct){
```

### Making it harder still

Here are some other ideas of ways in which you can beef up your captchas...

If you have a UI that lets you drag things around, you can display a list of five random numbers in a table, one number per row. To solve it, all you have to do is drag them into order with the lowest first and highest last. Or maybe in descending order.

Instead of a purely text question, draw a number of shapes (circle, square, triangle, star and so on) in a

rectangle and fill them in with different colours. Make sure the shapes do not overlap horizontally and have maybe 10-15 shapes. Then the captcha question can be something like these:

- What colour is the third circle from the left?
- What is the shape and colour of the shape to the left of the blue star?

For bonus points, draw fewer but bigger shapes, and put some shapes inside coloured rectangles.

- What is the shape and colour of the second shape inside the red rectangle? **Answer** yellow circle.
- What is the colour and shape of the third shape inside the blue rectangle? **Answer** blue circle.

You can turn the difficulty up a notch by putting numbers inside some of the shapes and then asking questions like:

- What colour is the rectangle containing a circle with the number 54 in it? **Answer** red.
- Which colour is the rectangle that contains no numbers? **Answer** black.
- What number is in the red circle left of a blue star? **Answer** 23.

Creating a captcha can be fun and, with a bit of imagination, you can create something that will take quite a lot of programming to defeat. Most websites are probably not important enough to warrant that. If yours is and your captcha proves difficult, you might provide extra employment to individuals who solve captchas for 1p each in distant countries.

The files, which are PHP plus an HTML file, are on GitHub – the link to the files is <https://github.com/David-H-Bolton/Projects/>. **LXF**

Showing the hashing of an answer and its salt.

Answer:	Five
Salt:	XYZ1234
Combined:	FiveXYZ1234
Hashed:	51e3ebbe0fed13bb38d2106b347e596287d7d910bc60119c2870a9a58b6eed7d

» **WE DON'T MIND IF YOU'RE HUMAN** Subscribe now at <http://bit.ly/LinuxFormat>

NEXT MONTH

**LXF323**  
will be on sale  
Tuesday 10th  
December  
2024

# CHRISTMAS Pi-JECTS

It's the most wonderful time of the year to be a maker! Build fun, cool and inspirational projects.

## Get developing!

Kick off the new year with a better integrated development environment. We test the best open source options...

## Better open source maps

Get out into the real world and start contributing to one of the best open source projects: OpenStreetMap.

## Protect your links

We explore Linkwarden, the tool for creating your own safely archived and shared web links and pages.

## Need input!

We explore how to improve your input life with enhanced keyboards that can boost your tippy-tappy typing!

Content of future issues subject to change. We might have had one too many sherries...

CREDIT: Magictorch

# LINUX FORMAT

The #1 open source mag

Future Publishing Limited, Quay House, The Ambury, Bath, BA11 1UA  
Email [contact@linuxformat.com](mailto:contact@linuxformat.com)

### EDITORIAL

Editor-in-chief Neil Mohr

Art editor Fraser McDermott

Production editor Katharine Davies

Content director Marc Chacksfield

Group art director Warren Brown

### Editorial contributors

Mike Bedford, Denise Bertacchi, Jonni Bidwell, David Bolton, Neil Bothwick, Stuart Burns, Ferenc Deák, Nate Drake, Brandon Hill, Zhiye Liu, Christopher Livingston, Jon Masters, Nick Peers, Les Pounder, Michael Reed, Mayank Sharma, Shashank Sharma, Chris Szweczyk

### Cover illustration

Linux Fedora and the Infinity design logo are trademarks of Red Hat, Inc.

Ubuntu is a trademark of Canonical Limited.

We are not endorsed by or affiliated with Canonical Limited or the Ubuntu project.

Raspberry Pi is a trademark of the Raspberry Pi Foundation.

Tux credit: Larry Ewing ([lewing@isc.tamu.edu](mailto:lewing@isc.tamu.edu)).

Content production Adequate Media Limited

### ADVERTISING

Commercial sales director Clare Dove

[clare.dove@futurenet.com](mailto:clare.dove@futurenet.com)

Advertising director Lara Jaggon

[lara.jaggon@futurenet.com](mailto:lara.jaggon@futurenet.com)

Account director Andrew Tilbury

[andrew.tilbury@futurenet.com](mailto:andrew.tilbury@futurenet.com)

### INTERNATIONAL LICENSING

Head of print licensing Rachel Shaw

Linux Format is available for licensing and syndication.

To find our more contact us at [licensing@futurenet.com](mailto:licensing@futurenet.com)

or view our content at [www.futurecontenthub.com](http://www.futurecontenthub.com)

### NEW SUBSCRIPTIONS & PAST ISSUES

Web [www.magazinesdirect.com](http://www.magazinesdirect.com)

### EXISTING SUBSCRIPTIONS

Web [www.mymagazine.co.uk](http://www.mymagazine.co.uk)

Subscription delays: Please allow up to seven days before contacting us about a late delivery to [help@magazinesdirect.com](mailto:help@magazinesdirect.com)

### MANAGE YOUR SUBSCRIPTION ONLINE WITH MYMAGAZINE

Visit [www.mymagazine.co.uk/FAQ](http://www.mymagazine.co.uk/FAQ) to view frequently asked questions or log in at [www.mymagazine.co.uk](http://www.mymagazine.co.uk)

### CIRCULATION

Newtrade & retail category director Ben Oakden

### PRODUCTION AND DISTRIBUTION

Group head of production Mark Constance

Production manager Nola Cokely

Senior ad production manager Jo Crosby

Digital editions manager Jason Hudson

### THE MANAGEMENT

Managing director technology group Paul Newman

Global head of design Rodney Dive

Commercial finance director Tania Brunning

Printed by William Gibbons & Sons

Distributed by Marketforce UK

121-141 Westbourne Terrace, London W2 6JR. [www.marketforce.co.uk](http://www.marketforce.co.uk)

For enquiries email: [mfcommunications@futurenet.com](mailto:mfcommunications@futurenet.com)

**Order and access past issues:** If you are an active subscriber, you have instant access to past issues through your iOS or Android device/s. Your digital magazine entitlement is available at no additional cost and no further action is required. Pocketmags library may not have access to the full archive of digital past issues. You will only be able to access the digital back issues as long as you are an active subscriber.

**To purchase single past issues (print format only):** Visit [www.magazinesdirect.com](http://www.magazinesdirect.com) (click on Single Issues tab) or email [help@magazinesdirect.com](mailto:help@magazinesdirect.com). Magazinesdirect.com is owned and operated by Future Publishing Limited.

**BAR rates for Linux Format:** £84.37 for UK, £171 for Europe, \$194 for USA, £149 for rest of world.

**Linux®** is the registered trademark of Linux Torvalds in the US and other countries. GNU/Linux is abbreviated to Linux throughout for brevity. Where applicable, code printed in this magazine is licensed under the GNU GPL v2 or later. See [www.gnu.org/copyleft.html](http://www.gnu.org/copyleft.html). All copyrights and trademarks are recognised and respected.

**Disclaimer:** All contents © 2024 Future Publishing Limited or published under licence. All rights reserved. No part of this magazine may be used, stored, transmitted or reproduced in any way without the prior written permission of the publisher. Future Publishing Limited (company number 2008885) is registered in England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All information contained in this publication is for information only and is, as far as we are aware, correct at the time of going to press. Future Publishing Limited cannot accept any responsibility for errors or inaccuracies in such information. You are advised to contact manufacturers and retailers directly with regard to the price of products/services referred to in this publication. Apps and websites mentioned in this publication are not under our control. We are not responsible for their contents or any other changes or updates to them. This magazine is fully independent and not affiliated in any way with the companies mentioned herein.

If you submit material to us, you warrant that you own the material and/or have the necessary rights/permissions to supply the material and you automatically grant Future Publishing Limited and its licensees a licence to publish your submission in whole or in part in any/all issues and/or editions of publications, in any format published worldwide and on associated websites, social media channels and associated products. Any material you submit is sent at your own risk and, although every care is taken, neither Future nor its employees, agents, subcontractors or licensees shall be liable for loss or damage. We assume all unsolicited material is for publication unless otherwise stated, and reserve the right to edit, amend and adapt all submissions. All contents in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your systems, peripherals or software through the use of any guide.

**We are committed to** only using magazine paper derived from responsibly managed, certified forestry and chlorine-free manufacture. The paper in this magazine was sourced and produced from sustainable managed forests, conforming to strict environmental and socioeconomic standards.

Future is an award-winning international media group and leading digital business. We reach more than 57 million international consumers a month and create world-class content and advertising solutions for passionate consumers online, on tablet & smartphone and in print.



**FUTURE**

Connectors.  
Creators.  
Experience  
Makers.


Future plc is a public  
company quoted on the  
London Stock Exchange  
(symbol: FUTR)  
[www.futureplc.com](http://www.futureplc.com)

Chief Executive Officer **Jon Steinberg**  
Non-Executive Chairman **Richard Huntingford**  
Chief Financial Officer **Sharjeel Suleman**

Tel: +44 (0)1225 442 244







The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

**EFF.ORG**


**ELECTRONIC FRONTIER FOUNDATION**

Protecting Rights and Promoting Freedom on the Electronic Frontier



# SAVVY SAVERS, STEP THIS WAY.

Discover cover options for car, home and travel insurance, plus broadband and more.



**GO.**  
**COMPARE**

Get more information and compare quotes at [Go.Compare](https://www.Go.Compare).